



Game Theory

through EXAMPLES

ERICH
PRISNER



MAA

CLASSROOM RESOURCE MATERIALS

VISIT...

LANZAROTE
Caliente.COM

Game Theory Through Examples

© 2014 by the Mathematical Association of America, Inc.

Electronic edition ISBN 978-1-61444-115-1

Game Theory Through Examples

Erich Prisner
Franklin University Switzerland



*Published and Distributed by
The Mathematical Association of America*

Council on Publications and Communications

Frank Farris, *Chair*

Committee on Books

Fernando Gouvêa, *Chair*

Classroom Resource Materials Editorial Board

Susan G Staples, *Editor*

Michael Bardzell

Jennifer Bergner

Caren L Diefenderfer

Christopher Hallstrom

Cynthia J Huffman

Paul R Klingsberg

Brian Lins

Mary Morley

Philip P Mummert

Barbara E Reynolds

Darryl Yong

CLASSROOM RESOURCE MATERIALS

Classroom Resource Materials is intended to provide supplementary classroom material for students—laboratory exercises, projects, historical information, textbooks with unusual approaches for presenting mathematical ideas, career information, etc.

101 Careers in Mathematics, 3rd edition edited by Andrew Sterrett

Archimedes: What Did He Do Besides Cry Eureka?, Sherman Stein

Calculus: An Active Approach with Projects, Stephen Hilbert, Diane Driscoll Schwartz, Stan Seltzer, John Maceli, and Eric Robinson

Calculus Mysteries and Thrillers, R. Grant Woods

Conjecture and Proof, Miklós Laczkovich

Counterexamples in Calculus, Sergiy Klymchuk

Creative Mathematics, H. S. Wall

Environmental Mathematics in the Classroom, edited by B. A. Fusaro and P. C. Kenshaft

Excursions in Classical Analysis: Pathways to Advanced Problem Solving and Undergraduate Research, by Hongwei Chen

Explorations in Complex Analysis, Michael A. Brilleslyper, Michael J. Dorff, Jane M. McDougall, James S. Rolf, Lisbeth E. Schaubroeck, Richard L. Stankewitz, and Kenneth Stephenson

Exploratory Examples for Real Analysis, Joanne E. Snow and Kirk E. Weller

Exploring Advanced Euclidean Geometry with GeoGebra, Gerard A. Venema

Game Theory Through Examples, Erich Prisner

Geometry From Africa: Mathematical and Educational Explorations, Paulus Gerdes

Historical Modules for the Teaching and Learning of Mathematics (CD), edited by Victor Katz and Karen Dee Michalowicz

Identification Numbers and Check Digit Schemes, Joseph Kirtland

Interdisciplinary Lively Application Projects, edited by Chris Arney

Inverse Problems: Activities for Undergraduates, Charles W. Groetsch

Keeping it R.E.A.L.: Research Experiences for All Learners, Carla D. Martin and Anthony Tongen

Laboratory Experiences in Group Theory, Ellen Maycock Parker

Learn from the Masters, Frank Swetz, John Fauvel, Otto Bekken, Bengt Johansson, and Victor Katz

Math Made Visual: Creating Images for Understanding Mathematics, Claudi Alsina and Roger B. Nelsen

Mathematics Galore!: The First Five Years of the St. Marks Institute of Mathematics, James Tanton

Methods for Euclidean Geometry, Owen Byer, Felix Lazebnik, and Deirdre L. Smeltzer

Ordinary Differential Equations: A Brief Eclectic Tour, David A. Sánchez

Oval Track and Other Permutation Puzzles, John O. Kiltinen

Paradoxes and Sophisms in Calculus, Sergiy Klymchuk and Susan Staples

A Primer of Abstract Mathematics, Robert B. Ash

Proofs Without Words, Roger B. Nelsen

Proofs Without Words II, Roger B. Nelsen

Rediscovering Mathematics: You Do the Math, Shai Simonson

She Does Math!, edited by Marla Parker

Solve This: Math Activities for Students and Clubs, James S. Tanton

Student Manual for Mathematics for Business Decisions Part 1: Probability and Simulation, David Williamson, Marilou Mendel, Julie Tarr, and Deborah Yoklic

Student Manual for Mathematics for Business Decisions Part 2: Calculus and Optimization, David Williamson, Marilou Mendel, Julie Tarr, and Deborah Yoklic

Teaching Statistics Using Baseball, Jim Albert

Visual Group Theory, Nathan C. Carter

Which Numbers are Real?, Michael Henle

Writing Projects for Mathematics Courses: Crushed Clowns, Cars, and Coffee to Go, Annalisa Crannell, Gavin LaRose, Thomas Ratliff, and Elyn Rykken

MAA Service Center
P.O. Box 91112
Washington, DC 20090-1112
1-800-331-1MAA FAX: 1-301-206-9789

Contents

Preface	xvi
1 Theory 1: Introduction	1
1.1 What's a Game?	1
1.2 Game, Play, Move: Some Definitions	1
1.3 Classification of Games	2
Exercises	3
2 Theory 2: Simultaneous Games	4
2.1 Normal Form—Bimatrix Description	4
2.1.1 Two Players	4
2.1.2 Two Players, Zero-sum	5
2.1.3 Three or More Players	5
2.1.4 Symmetric Games	6
2.2 Which Option to Choose	6
2.2.1 Maximin Move and Security Level	6
2.2.2 Dominated Moves	7
2.2.3 Best Response	8
2.2.4 Nash Equilibria	9
2.3 Additional Topics	13
2.3.1 Best Response Digraphs	13
2.3.2 2-Player Zero-sum Symmetric Games	14
Exercises	15
Project 1: Reacting fast or slow	18
3 Example: Selecting a Class	19
3.1 Three Players, Two Classes	19
3.1.1 "I like you both"	19
3.1.2 Disliking the Rival	21
3.1.3 Outsider	21
3.2 Larger Cases	22
3.3 Assumptions	23
Exercises	23
Project 2	23
Project 3	24
Project 4	24

4	Example: Doctor Location Games	25
4.1	Doctor Location	25
4.1.1	An Example Graph	26
4.1.2	No (Pure) Nash Equilibrium?	27
4.1.3	How Good are the Nash Equilibria for the Public?	28
4.2	Trees	28
4.3	More than one Office (optional)	31
	Exercises	31
	Project 5: Doctor location on MOPs	33
	Project 6	33
	Project 7	33
5	Example: Restaurant Location Games	34
5.1	A First Graph	35
5.2	A Second Graph	36
5.3	Existence of Pure Nash Equilibria	37
5.4	More than one Restaurant (optional)	38
	Exercises	39
6	Using Excel	42
6.1	Spreadsheet Programs like Excel	42
6.2	Two-Person Simultaneous Games	43
6.3	Three-Person Simultaneous Games	43
	Exercises	43
	Project 8: Simultaneous Quatro-Uno	44
	Project 9: Restaurant Location Games	44
	Project 10: 5 Knights	45
	Project 11: 5 Cardinals	45
7	Example: Election I	47
7.1	First Example	47
7.2	Second Example	48
7.3	The General Model	50
7.4	Third Example	50
7.5	The Eight Cases	51
7.6	Voting Power Indices (optional)	51
	Exercises	52
8	Theory 3: Sequential Games I: Perfect Information and no Randomness	53
8.1	Extensive Form: Game Tree and Game Digraph	53
8.2	Analyzing the Game: Backward Induction	56
8.2.1	Finite Games	56
8.2.2	The Procedure	57
8.2.3	Zermelo’s Theorem	59
8.3	Additional Topics	59
8.3.1	Reality Check	59
8.3.2	Playing it Safe—Guaranteed Payoffs	61
8.3.3	Two-person Zero-sum Games	63

8.3.4	Breaking Ties	64
8.3.5	Existing Games	64
8.3.6	Greedy Strategies	65
	Exercises	65
	Project 12: TAKE SOME	69
	Project 13: WHO's NEXT(n)	69
	Project 14: LISA's GAME	69
	Project 15: 2-AUCTION	69
	Project 16: 3-AUCTION	69
9	Example: Dividing A Few Items I	70
9.1	Greedy Strategy	70
9.2	Backward Induction	71
9.2.1	Game Tree	71
9.2.2	Game Digraph	71
9.2.3	Example: Game Digraph for ABBAB	72
9.3	An Abbreviated Analysis	72
9.3.1	Why it Matters: Complexity (optional)	74
9.4	Bottom-Up Analysis	75
9.5	Interdependencies between the Items (optional)	76
	Exercises	76
10	Example: Shubik Auction I	77
	Exercises	79
	Project 17: SHUBIK AUCTION	79
11	Example: Sequential Doctor and Restaurant Location	80
11.1	General Observations for Symmetric Games	80
11.2	Doctor Location	81
11.3	Constant-Sum Games	82
11.4	Restaurant Location	83
11.5	Nash Equilibria and First Mover Advantage for Symmetric Games	84
	Exercises	84
	Project 18	85
	Project 19: Hostile versus Friendly Play	85
12	Theory 4: Probability	86
12.1	Terminology	86
12.2	Computing Probabilities	88
12.2.1	Equally Likely Simple Events	88
12.2.2	Simple Events not Equally Likely	88
12.3	Expected Value	89
12.4	Multistep Experiments	91
12.4.1	Probability Trees	91
12.4.2	Conditional Probabilities	91
12.4.3	Probability Digraphs	93
12.5	Randomness in Simultaneous Games	94
12.6	Counting without Counting	95

Exercises	95
Project 20: Tennis	97
Project 21: Final Exam	98
13 France 1654	99
Exercises	101
14 Example: DMA Soccer I	102
14.1 1-Round 2-Step Experiment for Given Player Distributions	103
14.2 Expected Goal Difference for the One-Round Game	104
14.3 3-Rounds Experiment for Given Player Distributions	105
14.4 Static Three-round Game	107
14.5 Static Nine-round DMA Soccer	108
Exercises	108
Project 22: DMA6* Soccer	109
Project 23: DMA7* Soccer	109
15 Example: Dividing A Few Items II	110
15.1 Goals of Fairness and Efficiency	110
15.1.1 Fairness	110
15.1.2 Efficiency	111
15.1.3 Three Additional Features	111
15.1.4 Mechanism Design	112
15.2 Some Games	112
15.2.1 Selecting one by one Games	112
15.2.2 Cut and Choose	112
15.2.3 Random and Exchange	113
15.3 Examples	113
15.4 Comparison of the Games for Seven Items and Complete Information	115
15.4.1 Opposing or Similar Preferences	116
15.5 Incomplete Information	118
Exercises	119
Project 24: Dividing five items A	120
Project 25: Dividing five items B	120
16 Theory 5: Sequential Games with Randomness	121
16.1 Extensive Form Extended	121
16.2 Analyzing the Game: Backward Induction again	121
16.3 Decision Theory: Alone against Nature	122
Exercises	125
Project 26: Job Interviews	127
Project 27: 5 Envelopes	127
Project 28: Oh-No or Oh-No6	128
Project 29: 3×4 version of Polyomino REC THE SQUARE with randomness	128
17 Example: Sequential Quiz Show I	129
17.1 Candidates with Little Knowledge	129
17.1.1 More May be Less	130

17.2 One Candidate Knows More	131
17.2.1 Cindy Knows one Answer to be False	133
Exercises	133
Project 30: SEQUENTIAL QUIZ SHOW, clever Ann	134
Project 31: SEQUENTIAL QUIZ SHOW, clever Beth	134
18 Las Vegas 1962	135
Exercises	138
19 Example: Mini Blackjack and Card Counting	139
19.1 The Basic Game	139
19.2 Playing against the House	142
19.2.1 How Likely are the Distributions?	143
19.2.2 Betting High and Low	145
19.2.3 Reshuffling	146
Exercises	147
Project 32: 8 ROUNDS BLACK OR WHITE	147
Project 33: x ROUNDS RED, GREEN, OR BLUE	148
Project 34: MINI BLACKJACK	148
20 Example: Duel	149
20.1 One Bullet	149
20.1.1 Analysis of One-bullet Variants with Increasing Probabilities without Computer Help	150
20.1.2 Analysis of DUEL(1, 1 m, 2m, 3m, ...)	151
20.2 Two or more Bullets	152
20.2.1 A few Cases of DUEL(2, 2 m, 2m, 3m, ...)	153
Exercises	154
Project 35: Drunk Adam	154
Project 36: How more dangerous weapons affect the state budget and the health of citizens	154
Project 37: Selecting m between 0.04 and 0.13	154
Project 38: What duels are best for society?	155
21 Santa Monica in the 50s	156
22 Theory 6: Extensive Form of General Games	159
22.1 Extensive Form and Information Sets	159
22.2 No Backward Induction for Imperfect Information	163
22.3 Subgames	164
22.4 Multi-round Games	165
22.5 Why Trees for Imperfect Information?	165
Exercises	166
23 Example: Shubik Auction II	169
23.1 Possible Sudden End	169
23.2 Imperfect and Incomplete Information	172
23.3 The Auctioneer Enters the Game (optional)	172
Exercises	174
Project 39	174

Project 40	174
Project 41: SHUBIK AUCTION(45, 35, 6, p)	175
Project 42: SHUBIK AUCTION(A, B, C, n, p)	175
24 Theory 7: Normal Form and Strategies	176
24.1 Pure Strategies	176
24.1.1 Reduced Pure Strategies	177
24.2 Normal Form	177
24.3 Using Tools from Simultaneous Games for the Normal Form	180
24.4 Subgame Perfectness	180
24.5 Special Case of Sequential Games with Perfect Information	182
Exercises	182
25 Example: VNM POKER and KUHN POKER	185
25.1 Description	185
25.2 VNM POKER	187
25.3 KUHN POKER	190
Exercises	191
26 Example: Waiting for Mr. Perfect	193
26.1 The Last Round	193
26.2 The Eight Pure Strategies	194
26.3 Computing the Payoffs	194
26.4 Domination	195
26.5 The Reduced Normal Forms in the Three Cases	196
26.5.1 The Case $p_2 + 2p_3 < 1$	196
26.5.2 The Case $p_2 + 2p_3 > 1$	197
26.5.3 The Case $p_2 + 2p_3 = 1$	197
Project 43	198
Project 44	198
Project 45	198
Project 46	198
Project 47	198
Project 48	198
27 Theory 8: Mixed Strategies	199
27.1 Mixed Strategies	199
27.1.1 Best Response	200
27.1.2 Brown's Fictitious Play	200
27.1.3 Mixed Maximin Strategy, Mixed Security Level, and Linear Programs	202
27.2 Mixed Nash Equilibria	203
27.2.1 Two-player Zero-sum Games	203
27.2.2 Non-Zero-sum Games	203
27.3 Computing Mixed Nash Equilibria	204
27.3.1 Small Two-player Zero-sum Games (optional)	205
27.3.2 Solving Small non Zero-sum Two-player Games by Solving Equations (optional)	207
Exercises	209
Project 49: Balanced 3-spinner duel with five options	210

Project 50: Balanced 3-spinner duel	210
Project 51: COLONEL BLOTTO(4, 9, 9)	210
Project 52: Iterated COLONEL BLOTTO	211
Project 53: Simultaneous Quatro-Uno	211
Project 54	211
Project 55: 4-round Waiting for Mr. Perfect	211
28 Princeton in 1950	212
29 Example: Airport Shuttle	215
29.1 The Simple Model	215
29.1.1 To the Airport	215
29.1.2 From the Airport	218
29.1.3 Combining Both	218
29.2 Impatient Customers	219
Exercises	219
30 Example: Election II	220
30.1 Left Over from Election I	220
30.2 More Effort into Large Districts	221
30.3 Defend Where Ahead or Attack Where Weak?	222
30.4 Is Larger Better?	223
30.5 ELECTION(7, 8, 13 − 1, −1, 2 x, x)	223
Exercises	224
31 Example: VNM POKER(2, r, m, n)	225
31.1 The Case $\frac{n}{m} \geq 2 - \frac{1}{r}$	226
31.2 Best Responses	227
31.3 Reading the Opponent (optional)	227
31.4 Mixed Nash Equilibrium for $\frac{n}{m} \leq 2 - \frac{1}{r}$	228
31.5 Small Changes in the Parameters	229
Exercises	230
32 Theory 9: Behavioral Strategies	231
32.1 Behavioral versus Mixed Strategies	232
32.1.1 Calculating Mixed Strategies from Behavioral Strategies	233
32.1.2 Calculating Behavioral Strategies from Mixed Strategies for a Game Tree with Perfect Recall	233
32.1.3 Kuhn’s Theorem	235
Exercises	235
33 Example: Multiple-Round Chicken	237
33.1 Ordinary Chicken	237
33.2 Two-round Chicken	238
33.2.1 Generalized Backward Induction, using the Extensive Form	238
33.2.2 Working with the Normal Form	240
33.2.3 Connections between the two Approaches	240
33.3 Three-round Chicken	241

Exercises	243
Project 56	243
Project 57	243
Project 58	243
34 Example: DMA Soccer II	244
34.1 Multi-round Simultaneous Games	244
34.2 Information Sets and Moves	245
34.3 The Optimal Third Move in Selected Cases	246
34.3.1 A Detailed Example: (2, 2) versus (3, 1)	246
34.3.2 A Second Example: (1, 3) versus (2, 2)	249
34.4 The Optimal Second Move for Seven Positions	249
34.5 Couldn't We Analyze the Whole Game?	251
34.6 How Good a Model is it?	251
35 Example: Sequential Quiz Show II	252
35.1 Fixed Coalitions	252
35.1.1 Ann and Cindy Form a Coalition	252
35.1.2 Ann and Beth Form a Coalition	254
35.1.3 Beth and Cindy Form a Coalition	254
35.2 Which Coalition Will Form?	254
35.2.1 Fixed 50:50 Split	254
35.3 Another Variant: Split can be Negotiated	255
35.4 The Grand Coalition	256
35.4.1 The Core	256
35.4.2 The Shapley Value	256
Exercises	257
Project 59	258
36 Example: VNM POKER(4, 4, 3, 5)	259
36.1 Mixed Nash Equilibria	259
36.2 Performance of Pure Strategies against the Mixed Nash Equilibria	261
37 Example: KUHN POKER(3, 4, 2, 3)	264
37.1 From Behavioral Strategies to Mixed Strategies to Expectations	265
37.2 From Mixed Strategies to Behavioral Strategies	266
Exercises	266
38 Example: End-of-Semester Poker Tournament	268
38.1 Expectations	269
38.2 Odds	270
38.2.1 Many Rounds	273
38.3 The Favorite in Knockout Tournaments	273
38.4 Similarity of the DNA (optional)	274
38.5 How to Create your own Tournament	274
Exercises	275
Project 60	275
Project 61	276

Project 62	276
39 Stockholm 1994	277
Bibliography	280
Index	284

Preface

Welcome to game theory, the mathematical theory of how to analyze games and how to play them optimally. Although “game” usually implies fun and leisure, game theory is a serious branch of mathematics. Games like blackjack, poker, and chess are obvious examples, but there are many other situations that can be formulated as games. Whenever rational people must make decisions within a framework of strict and known rules, and where each player gets a payoff based on the decisions of all the players, we have a game. Examples include auctions, negotiations between countries, and military tactics. The theory was initiated by mathematicians in the first half of the last century, but since then much research in game theory has been done outside of mathematics.

This book gives an introduction to game theory, on an elementary level. It does not cover all areas of a field that has burgeoned in the last sixty years. It tries to explain the basics thoroughly, so that it is understandable for undergraduates even outside of mathematics, I hope in an entertaining way. The book differs from other texts because it emphasizes examples. The theory is explained in nine theory chapters, and how it is applied is illustrated in twenty-four example chapters, where examples are analyzed in detail. Each example chapter uses tools developed in theory chapters.

Audience

The text can be used for different purposes and audiences. It can be used by students of economics, political science, or computer science. Undergraduate mathematics students should enjoy the book as well and profit from its approach. The book may also be used as a secondary text, or for independent study, as its many concrete examples complement expositions of the theory.

I think there is another audience that could profit from the book, the one I had in mind when I wrote the book, which is undergraduates who take a mathematics class to fulfill general education or quantitative reasoning requirements. The text began as an online text for a first year seminar at Franklin College in 2007, when I couldn’t find an appropriate text. Since then, I have used versions of the text in our first year seminar and in a liberal arts course I have called Introduction to Game Theory. Game theory and my approach to it are well suited for these courses because:

- The underlying mathematics is basic: for example, finding minima and maxima, computing weighted averages and working systematically through trees or digraphs.
- The mathematical topics that are covered—probability, trees, digraphs, matrices, and algorithms—are among the most useful tools used in applications.
- Game theory is applied mathematics. Some would claim it is not even mathematics, but part of economics. The “why are we doing this?” question has an obvious answer. Students accept the fact that mathematical tools have to be developed for studying and analyzing games.
- Game theory gives opportunities for students to do projects.
- Game theory allows a playful approach.

Features of the book

There are thirty-eight chapters of three types:

- Nine chapters present the basic theory. They are listed as theory chapters in the table of contents. In most courses all these chapters would be covered. There is an additional chapter explaining how to use the Excel files, which is also required reading.
- Five chapters describe notable landmarks in the history and development of game theory: the 1654 letters between Pascal and Fermat, which mark the beginning of probability theory; Princeton University before 1950, where game theory started as an academic discipline; RAND Corporation in the 1950s and the early optimism there about what game theory could achieve; casino games; and the Nobel prizes awarded to game theorists. The chapters provide background about some of the important persons in the development of game theory, and discuss game theory's role in society. The chapters might be for reading at home for students. They could also be vehicles for class discussions.
- The real core of the manuscript are the twenty-four chapters that treat concrete examples. They distinguish the book from others. The chapters use the theory developed in the theory chapters, and they usually build on at most one other example chapter, so an instructor can select those that seem most appropriate. In my classes, I usually discuss about eight examples. Some of the example chapters provide glimpses into areas that are not covered in the theory chapters. Examples of this are trees in Chapter 4, voting power indices in Chapter 7, complexity and binomial coefficients in Chapter 9, statistics, mechanism design, and incomplete information in Chapter 15, incomplete versus imperfect information in Chapter 23, work with parameters and algebra in Chapter 31, and cooperative games in Chapter 35.

There are two other main features of the book: the Excel spreadsheets and Javascript applets.

- The Javascript applets are small games, in which students can try out many of the games discussed and analyzed in the examples. Students can play against the computer or against students. Before reading the analysis students should have played the games and developed some ideas on how to play the game. The Javascripts are also used to confirm theoretical results.
- The Excel spreadsheets either are generally usable for any simultaneous 2-player game or for every 2-player game in normal form, or they are designed for specific games. They give the students tools to check what is described in the text without having to do many routine or tedious calculations, and also tools to apply “what if” analysis to games. Students will see, sometimes in unexpected ways, how changing parameters may change strategies, outcomes, and payoffs.

In my opinion Excel is the best tool for the game theory in this book, better than Mathematica or Maple or MathCad. Since all steps are elementary (no functions are used except, “Min”, “Max”, “Sum”, and “If”), everything could also be done by the student on a sheet of paper, at least in principle. Every student should learn some Excel anyway, another feature which makes a course based on this book very well-suited for general education core requirements.

For the Instructor

The book's presentation is informal. It emphasizes not formulas but understanding, interpretation, and applicability. In my opinion, formalism and proof are perfect tools and unavoidable for professional mathematicians but are of limited value for students, even to some extent for undergraduate students of mathematics. Some proofs of theorems are given, for others informal reasonings, but often students have to rely on evidence from the examples and simulations given. I try to avoid using complicated formulas. When I present a formula, I try to explain it in words. It helps students to concentrate on concrete examples with concrete

numbers. When the book introduces parameters, the analysis carries them forward as a first step towards abstraction.

Although the book is elementary and avoids formalism, it may not be easy reading. Some books consider only 2 by 2 bimatrix games or simple real-life situations, but the examples in this book are sometimes complex. In my opinion, the power of mathematics cannot be appreciated by looking only at small examples. In them little mathematics is needed and common sense shows the best strategy. The students who use this book will use tables of numbers with ten to twenty rows and columns. Although this requires more effort than smaller examples, the Excel sheets will help with the calculations, and in the end the results should justify the extra effort. Students will see that mathematical theory sometimes produces unexpected results.

Though the examples are sometimes complex, they are still far removed from most real games or from serious modeling. I try to make this clear throughout the book. That does not mean that no lessons can be drawn from them!

It is common knowledge that mathematics can be learned only by doing. The book contains sixty or so projects that ask for the solution of a game and the presentation of it. Experimenting and going beyond what is described in the text is in my opinion crucial for mastering the material. Some projects are difficult and some are open-ended. I tell my students to do as much as they can, and to do that well!

Figure 1 shows how the chapters are related. The eight theory chapters are squares, the Excel is a diamond, the five history chapters are stars, and the example chapters are circles. The backbone of every course would be the theory chapters and the Excel chapter, chapters 1, 2, 6, 8, 12, 16, 22, 24, 27, and 32. They should in my opinion be covered in every course on game theory. Some of the history chapters, chapters 13, 18, 21, 28, and 39, could be assigned as reading and material for class discussion.




Figure 1. Structure of the chapters

Each example chapter is attached by a line to the earliest theory chapter after which it can be covered. Some of the example chapters form pairs, like Election and Election II, which are enclosed by dashed lines. I would not recommend covering a part II chapter without having covered part I. There are two larger groups, the location games of chapters 4, 5, and 11, and the poker chapters 25, 31, 36, 37, 38. In these cases I would recommend at least one of chapters 4 or 5 before chapter 11, and covering chapter 25 before covering any of chapters 31, 36, 37, or 38.

As more tools become available, the example chapters become more complex and difficult, so later chapters are generally more difficult than early ones. In Figure 1 I indicate the difficulty of the example chapters by showing them above or below the theory backbone. The example chapters below the backbone I consider

to be easier, those above more difficult. But sometimes it is possible to skip the difficult parts. For example, chapters 4, 5, and 11 all contain proofs that relate games to the structure of graphs, and this may be of interest to mathematics or computer science majors. But the games provide all readers with instructive examples of simultaneous and sequential games.

When I teach my course, I usually cover at least one example chapter after each theory chapter. For a general education mathematics course an appropriate set may be chapters 3, 4 (just Section 4.1), 7, 10, 17, 23, 26, 30, and 35.

Though the poker chapters 31, 36, and 37, are complex, I usually cover one of them in my course, since I end the semester with a robot poker tournament similar to the one described in chapter 38. The tournament gets the student's attention and is usually a lot of fun! Instructors can modify the applets to run their own tournaments and even vary the rules of the game.

Acknowledgements

Early versions of this book were used as an online text for seven classes at Franklin College (now Franklin University) between 2007 and 2013. I want to thank all my students for fighting their way through these often incomplete and imperfect pages, and for the valuable feedback they gave me. I want to thank in particular Samy Al-Ansari, Nurbek Almakuchukov, Meghan Canale, Eric Gors, Ruth Kappes, Matej Kristic, Grigore Tocitu, Karam Zaidkilani, Conner Crawford, Chikara Davis, James Jasper, Artyom Machekin, Nicolaus Owens, Megan Austin, Dylan Bontempo, Maija Butler, Andriana Friel, Rene Musech, Sannan Panaruna, Aaron Brinckerhoff, Caitlin Curtis, Mihael Djambazov, Mariah Grubb, Sophie Winckel, Nathan Ayers, Jen-nalee Bauman, Alex Bromwell, Stephanie Erichsen, Sachint Goel, Segev Goldberg, Andrea Hak, Michael Radomile, Cooper Stoulil, Jennifer Byram, Arnulfo Hermes, Connor McCormack, Benjamin Smick, Paige Stumbough, Emil Akhmetov, (Kevin) Yikfai Chan, Todd Gilbert, Ishan Hill, Anna Hixson, Adrian Lewis, Ariane Mottale, Lily Rybarczyk, Adam Taylor, and Senya Webster.

I also want to express my gratitude to Franklin University Switzerland for granting me released time during the years I have worked on writing this book.

I am grateful to Jerry Bryce and the Classroom Resource Materials Board of Editors, and would thank the following members for reviewing my manuscript: Phil Mummert, Phil Straffin, Susan Staples, Holly Zullo, Sister Barbara Reynolds, Cynthia Woodburn, and Diane Herrmann. Their remarks greatly helped to improve the book.

But my very deepest thanks go to the CRM Editor Jerry Bryce for going very carefully, sentence by sentence, over the text and correcting typos, bad English, and pointing out the many cases of unclearness in earlier versions of the manuscript. Any errors that remain are my sole responsibility. Without Jerry, the book would not have been possible.

Many thanks also to Underwood Dudley for copy editing the manuscript very carefully and greatly improving the manuscript.

I also want to thank the MAA acquisitions editors Don Albers, Zaven Karian, and Steve Kennedy, and the whole MAA editing and production team, in particular Carol Baxter and Beverly Ruedi, for help with the production process.

CHAPTER 1

Theory 1: Introduction

1.1 What's a Game?

Every child understands what games are. When someone overreacts, we sometimes say “it’s just a game.” Games are often not serious. *Mathematical* games, which are the subject of this book, are different. It was the purpose of game theory from its beginnings in 1928 to be applied to serious situations in economics, politics, business, and other areas. Even war can be analyzed by mathematical game theory. Let us describe the ingredients of a mathematical game.

Rules Mathematical games have strict rules. They specify what is allowed and what isn’t. Though many real-world games allow for discovering new moves or ways to act, games that can be analyzed mathematically have a rigid set of possible moves, usually all known in advance.

Outcomes and payoffs Children (and grown-ups too) play games for hours for fun. Mathematical games may have many possible **outcomes**, each producing **payoffs** for the players. The payoffs may be monetary, or they may express satisfaction. You want to win the game.

Uncertainty of the Outcome A mathematical game is “thrilling” in that its outcome cannot be predicted in advance. Since its rules are fixed, this implies that a game must either contain some random elements or have more than one player.

Decision making A game with no decisions might be boring, at least for the mind. Running a 100 meter race does not require mathematical skills, only fast legs. However, most sport games also involve decisions, and can therefore at least partly be analyzed by game theory.

No cheating In real-life games cheating is possible. Cheating means not playing by the rules. If, when your chess opponent is distracted, you take your queen and put it on a better square, you are cheating, as in poker, when you exchange an 8 in your hand with an ace in your sleeve. Game theory doesn’t even acknowledge the existence of cheating. We will learn how to win without cheating.

1.2 Game, Play, Move: Some Definitions

The complete set of rules describes a **game**. A **play** is an instance of the game. In certain situations, called **positions**, a player has to make a decision, called a **move** or an **action**. This is not the same as strategy. A **strategy** is a plan that tells the player what move to choose in every possible position.

Rational behavior is usually assumed for all players. That is, players have preferences, beliefs about the world (including the other players), and try to optimize their individual payoffs. Moreover, players are aware that other players are trying to optimize their payoffs.

1.3 Classification of Games

Games can be categorized according to several criteria:

- How many players are there in the game? Usually there should be more than one player. However, you can play roulette alone—the casino doesn’t count as player since it doesn’t make any decisions. It collects or gives out money. Most books on game theory do not treat one-player games, but I will allow them provided they contain elements of randomness.
- Is play simultaneous or sequential? In a **simultaneous game**, each player has only one move, and all moves are made simultaneously. In a **sequential game**, no two players move at the same time, and players may have to move several times. There are games that are neither simultaneous nor sequential.
- Does the game have random moves? Games may contain random events that influence its outcome. They are called **random moves**.
- Do players have perfect information? A sequential game has **perfect information** if every player, when about to move, knows all previous moves.
- Do players have **complete information**? This means that all players know the structure of the game—the order in which the players move, all possible moves in each position, and the payoffs for all outcomes. Real-world games usually do not have complete information. In our games we assume complete information in most cases, since games of incomplete information are more difficult to analyze.
- Is the game zero-sum? **Zero-sum games** have the property that the sum of the payoffs to the players equals zero. A player can have a positive payoff only if another has a negative payoff. Poker and chess are examples of zero-sum games. Real-world games are rarely zero-sum.
- Is communication permitted? Sometimes communication between the players is allowed before the game starts and between the moves and sometimes it is not.
- Is the game cooperative or non-cooperative? Even if players negotiate, the question is whether the results of the negotiations can be enforced. If not, a player can always move differently from what was promised in the negotiation. Then the communication is called “cheap talk”. A **cooperative game** is one where the results of the negotiations can be put into a contract and be enforced. There must also be a way of distributing the payoff among the members of the coalition. I treat cooperative games in Chapter 35.

Student Activity Play ten rounds in the applets for each one of the games [LisaGame](#), [QuatroUno](#), and [Auction](#). In the first two you can play against a (well-playing) computer, but in the third the computer serves only as auctioneer and you need to find another human player. Categorize each game as simultaneous or sequential or neither, and determine whether randomness is involved, whether the game has perfect information, and whether it is zero-sum. Determine the number of players in each game, and justify your conclusions.

Modeling Note Analyzing games like parlor games or casino games may seem to be enough motivation to develop a theory of games. However, game theory has higher aims. It provides tools that can be applied in many situations where two or more persons make decisions influencing each other.

A **model** is an abstract, often mathematical, version of reality. In this book a model is a game, which is supposed to yield some insight into a real-world situation. It is important not to confuse the model with reality—in reality there are almost never totally strict rules and players almost always have more options than they think they have, more than what the model allows.

In this book we also will try to model some real-world situations as games, but the approach taken is cautious. Whenever we try to model a real-life situation, we will

discuss in detail the assumptions of the model and whether the conclusions from the model are relevant. Whether game theory can be useful in real life is something for each reader to decide.

Exercises

1. In English auction, an item is auctioned. People increase bids in increments of \$10, and the player giving the highest bid gets the item for that amount of money. Give reasons why the auctioneer would be considered a player of the game, or reasons why he or she would not. Does the game contain random moves? Is it zero-sum? Would a real-world art auction have complete information?
2. In roulette, would the croupier be considered to be a player? Does the game contain random moves? Is it zero-sum? Can players increase their chances of winning if they form a coalition and discuss how to play before each round?
3. In the well-known game rock, scissors, paper game, how many players are there? Is it simultaneous, or sequential, or neither, and, if it is sequential, does it have perfect information?
4. For poker, discuss number of players, whether it is sequential or simultaneous, or neither, and if it is sequential, whether it has perfect information. Discuss whether there are random moves. Is communication allowed in poker?
5. For blackjack discuss its number of players; whether it is sequential or simultaneous, or neither; and if it is sequential, whether it has perfect information. Discuss whether there are random moves. Is communication allowed in blackjack?
6. It's late afternoon and you are in a train traveling along a coastline. From time to time the train stops in villages, some of them nice, some of them ugly, and you can evaluate the niceness of the village immediately. The benefit of an evening and night spent at that village depends only on its niceness. You want to get off at the nicest village. Unfortunately you don't know how many villages are still to come, and you know nothing about how villages in this country normally look. Worse, you are not able to ask anybody, since you don't speak the language of the country. You also know that some (unknown) time in the evening the train will reach its destination where you will have to stay whether it is nice or not. Explain the features of this game, with emphasis on the informational issues. How would you play it? Give a reason for your strategy. Comment on whether we have complete or incomplete information here, and why.
(Initially I formulated this example in terms of marriage in a society where divorce is impossible, but I saw that this is a different game. Could you give some arguments why?)
7. In this more realistic version of the game suppose that you know that the train will stop in ten villages before it reaches its destination. How would you play now? Comment on whether we have complete or incomplete information here, and justify your comment.

CHAPTER 2

Theory 2: Simultaneous Games

In his story “Jewish Poker” the writer Ephraim Kishon describes how a man called Ervinke convinces the narrator to play a game called Jewish Poker with him. “You think of a number, I also think of a number”, Ervinke explains. “Whoever thinks of a higher number wins. This sounds easy, but it has a hundred pitfalls.” Then they play. It takes the narrator some time until he realizes that it is better to let Ervinke tell his number first. [K1961] Obviously this is a game that is not fair unless both players play simultaneously.

In this chapter we will start our journey through game theory by considering games where each player moves only once, and moves are made simultaneously. The games can be described in a table (called the game’s **normal form**). Then we discuss approaches that allow the players to decide which move they will choose, culminating with the famous Nash equilibrium.

2.1 Normal Form—Bimatrix Description

Imagine you want to describe a simultaneous game. We know that each player has only one move, and that all moves are made simultaneously. What else do we need to say? First, we must stipulate the number of players in the game. Second, we must list for each player all possible moves. Different players may have different roles and may have different options for moves. We assume that each player has only finitely many options. Players simultaneously make their moves, determine the **outcome** of the game, and receive their payoffs. We need to describe the payoffs for each outcome.

How many outcomes are possible? Each combination of moves of the players generates a different outcome. If there are n players, and player 1 has k_1 possible moves, player 2 has k_2 possible moves, and so on, then there are $k_1 \times k_2 \times \dots \times k_n$ possible outcomes. For each, n numbers would describe the payoffs for player 1, player 2, and so on.

In games where each player has infinitely many options, we may use methods of calculus for functions with two variables, but such games are not discussed in this book. We describe simultaneous games with randomness in Chapter 12.

2.1.1 Two Players

Here is an example of a simultaneous 2-player game:

Example 1 ADVERTISING: Two companies share a market, in which they currently make \$5,000,000 each. Both need to determine whether they should advertise. For each company advertising costs \$2,000,000 and captures \$3,000,000 from the competitor provided the competitor doesn’t advertise. What should the companies do?

Let’s call the two companies A and B. If both don’t advertise, they get \$5,000,000 each. If both advertise, both lower their gain to \$3,000,000. If A advertises, but B doesn’t, A gets \$6,000,000 and B only \$2,000,000, and conversely if B advertises and A doesn’t. The payoff pattern is shown in the following table. The numbers are in millions of dollars. The rows correspond to the options of player A, and the columns correspond to the options of player B. The entries are payoff for A and payoff for B provided the corresponding options are chosen, separated by a comma.

	B advertises	B doesn’t advertise
A advertises	3, 3	6, 2
A doesn’t advertise	2, 6	5, 5

Whenever we have two players, we often name them Ann and Beth. Assume Ann has k_1 options and Beth has k_2 options. We want to display the different payoffs for Ann and Beth, depending on the different choices they have. Each of the $k_1 \cdot k_2$ outcomes has payoffs for Ann and Beth attached. Usually this is visualized in a table, the **normal form** with n rows, corresponding to Ann’s options, and m columns, corresponding to Beth’s options. Such a table is called a $n \times m$ bimatrix. The entries in the cells are payoffs for Ann and Beth, separated by a comma.

2.1.2 Two Players, Zero-sum

A game is called **zero-sum** if the sum of payoffs equals zero for any outcome. That means that the winnings of the winning players are paid by the losses of the losing players.

For zero-sum two-player games, the bimatrix representation of the game can be simplified: the payoff of the second player doesn’t have to be displayed, since it is the negative of the payoff of the first player.

Example 2 Assume we are playing ROCK-SCISSORS-PAPER for one dollar. Then the payoff matrix is

	Rock	Scissors	Paper
Rock	0	1	−1
Scissors	−1	0	1
Paper	1	−1	0

The first cell says “0”, which stands for “0, 0” a payoff of 0 for both players. The second cell entry of “1” should be read as “1, −1”, a payoff of 1 for Ann which has to be paid by Beth, therefore a payoff of −1 for Beth.

2.1.3 Three or More Players

If we have more than two players, we need another systematic way to generate the needed $k_1 \cdot k_2 \cdot \dots \cdot k_n$ cells corresponding to the different outcomes, into which we write the n payoffs for the n players. Here is an example:

Example 3 LEGISLATORS’ VOTE: Three legislators vote whether they allow themselves a raise in salary of \$2000 per year. Since voters are observing the vote, there is some loss of face for a legislator to vote for a raise. Let’s assume that the legislators estimate that loss of face is worth \$1000 per year. What happens if all three vote at the same time? (This game is a variant of the game described in [K2007]).

This is a simultaneous three-player game. It is best visualized with two matrices. Player A chooses the matrix, B chooses the row, and C chooses the column. The payoffs (in thousands of dollars) are

A votes for a raise			A votes against a raise		
	C votes for raise	C votes against it		C votes for raise	C votes against it
B votes for raise	1, 1, 1	1, 1, 2	B votes for raise	2, 1, 1	0, -1, 0
B votes against	1, 2, 1	-1, 0, 0	B votes against	0, 0, -1	0, 0, 0

2.1.4 Symmetric Games

All our examples so far are **symmetric**: All players have the same options, and if the two players interchange their moves, the payoffs are also interchanged. More formally, for a 2-player game, let m_1, m_2 be moves and let $a(m_1, m_2)$ and $b(m_1, m_2)$ be Ann’s and Beth’s payoffs if Ann plays m_1 and Beth plays m_2 . Then $a(m_1, m_2) = b(m_2, m_1)$ and $b(m_1, m_2) = a(m_2, m_1)$ for symmetric games. That means that the entries in the j s row and the i s column is obtained from the entries in the i s row and j s column by interchanging the payoffs. For symmetric 3-player games, $a(m_1, m_2, m_3) = b(m_2, m_1, m_3) = b(m_3, m_1, m_2) = c(m_2, m_3, m_1) = c(m_3, m_2, m_1)$, and so on. Symmetric games are fair by design, giving the same chances to every player.

2.2 Which Option to Choose

It is useful to describe simultaneous games by a bimatrix, but what players want is advice on how to play. Game theory should (and will in some cases) provide players with a mechanism to find which move is best.

The mechanisms would refer to the bimatrix only, the solution would be the same no matter whether we face a casino game or a war, provided the corresponding matrices are the same. The essence of the game lies in the numbers in the bimatrix.

Such mechanisms are the content of this section. We will give three or four of them. Like advice from well-meaning uncles, they all have a good and convincing point, but since they concentrate on different features of the game, they don’t always lead to the same conclusion. We will discuss them first separately before investigating the relations between them.

2.2.1 Maximin Move and Security Level

Some people always expect the worst. No matter what she plays, a player (let’s call her Ann) may assume that the other players will always respond with moves that minimize Ann’s payoff. This may be justified in a two-player zero-sum game if Ann is so predictable that the other player always anticipate her move. In other cases the belief borders on paranoia, since the other players will not be interested in harming Ann but instead want to maximize their payoffs. Still, pessimistic Ann will evaluate her strategies in light of the worst expected case. She would concentrate, for any of her options, on her smallest possible payoff. If she believes that this is what she would get, then Ann would choose the option with highest value. This value is called the **maximin value** or **security level**. The option Ann will play is called a **maximin move (strategy)**, since it maximizes the minimum possible payoff. Playing the maximin move, the player can guarantee a payoff of at least the maximin value, no matter how the others are playing. To choose the maximin move, the player doesn’t have to know the payoffs of the other players.

In the ADVERTISING example, company A may fear that company B will advertise too if A advertises, yielding a payoff of 3 for A. If company A does not advertise, the worst that could happen would be company

B advertising with payoff of 2 for A. Therefore company A would advertise to maximize the worst possible payoff.

In the **LEGISLATORS' VOTE** example, the worst that could happen if A votes for a raise is that both others vote against, leaving A with a payoff of -1000 . If A votes against a raise, in the worst case (actually in three of four cases) A gets a payoff of 0, which is more than in the other case. Therefore A would vote against a raise if using the maximin principle.

In a two-player game the first player, Ann, would look at the rows of the bimatrix and in each row highlight the cell with her lowest payoff. Then she would select the row with the highest number highlighted. In the same way, the second player, Beth, when playing the maximin strategy would mark in each column the cell with lowest payoff for Beth, and then select the column with the highest number marked.

How do we treat ties, if two or more rows have the same minimum payoff for Ann? Ann could choose one of these moves, or alternate randomly between such moves. The latter leads to mixed strategies that are covered in Chapter 27.

2.2.2 Dominated Moves

A move, $M1$, for Ann **strictly dominates** another $M2$, if $M1$ always results in a higher payoff for Ann than $M2$. A rational player would never play a move that is strictly dominated by another one. Domination doesn't tell what to play but rather what not to play. In the rare case where one of Ann's moves strictly dominates all her other moves, this would turn into positive advice to play the move dominating all other moves.

In the **ADVERTISING** example "advertising" strictly dominates "not advertising" for both companies. Therefore both companies will advertise, when applying this mechanism.

It is no coincidence that the advice given by the maximin mechanism and the advice given by the rule not to play strictly dominated moves are the same for this example. Actually a player's maximin move is never strictly dominated by any of her other moves.

Advice for players could go further than to disregard strictly dominated moves. In particular, if player Ann believes that other players would also obey this rule, then we may disregard all strictly dominated moves in the game, not only for Ann but for all other players. However, this assumption about the other players' behavior is not automatic. It assumes that all players are rational and clever or experienced enough. Under the assumption that all players accept this belief in the rationality and sophistication of all players, we know that all players reduce the game by eliminating all strictly dominated moves. Then, in the reduced game, strict domination may occur where it had not before, and the same round of eliminations could be done to reduce the game further. The process of repeatedly reducing the game, as well as its result, a game that cannot be reduced any further since there are no strictly dominated moves, is denoted by **IESD—iterated elimination of strictly dominated moves**. Except in cases where the IESD result is a game with just one option for Ann, IESD is a method of excluding moves rather than telling what move to choose.

Here is a successful application of the IESD procedure:

Example 4 TWO BARS: Each one of two bars charges its own price for a beer, either \$2, \$4, or \$5. The cost of obtaining and serving the beer can be neglected. It is expected that 6000 beers per month are drunk in a bar by tourists, who choose one of the two bars randomly, and 4000 beers per month are drunk by natives who go to the bar with the lowest price, and split evenly in case both bars offer the same price. What prices would the bars select? [S2009]

The game, as all games considered so far, is symmetric. Let me illustrate in one instance how to compute the payoffs. If bar A charges \$2 and bar B charges \$4, then all natives will choose bar A.

Therefore bar A will serve 4000 beers to the natives, and 3000 beers to tourists, serving 7000 beers in total, making $7000 \cdot 2 = 14000$ dollars. Bar B will only serve 3000 beers to tourists, making $3000 \cdot 4 = 12000$ dollars.

The payoff matrix, with values in thousands of dollars, is

	2	4	5
2	10, 10	14, 12	14, 15
4	12, 14	20, 20	28, 15
5	15, 14	15, 28	25, 25

For each bar, move “4” strictly dominates move “2”, therefore we could eliminate both moves “2” to get the reduced game:

	4	5
4	20, 20	28, 15
5	15, 28	25, 25

Now, but not before the elimination, move “4” strictly dominates move “5”. Therefore we eliminate these moves for both players as well and arrive at a game with only one option, “4”, for each player, and a payoff of \$ 20000 for each. Therefore both players will choose \$4 as the price of the beer.

A weaker condition is weak domination. Ann’s move **weakly dominates** another one of her moves if it yields at least the same payoff for Ann in all cases generated by combinations of moves of the other players, and in at least one case an even better payoff. So the weakly dominating move is never worse than the weakly dominated one, and sometimes it is better. The common wisdom is that **iterated elimination of weakly dominated moves, IEWD** is not something that should be performed automatically. Weakly dominated moves may still be played, in particular in cases where the weakness of the weakly dominated move appears in a combination with other player’s moves that are known not to be played by them. This opinion is also based on different behavior of Nash equilibria (discussed in Section 2.4) under IEWD and IESD.

2.2.3 Best Response

Assume you will play a one-round simultaneous game against your friend tomorrow. Your friend has been thinking about her move, arrives on a decision what move to play, and writes it on a piece of paper so as not to forget it. You get a look at this paper without your friend noticing it. The game thus changes from simultaneous to sequential with perfect information. The move you play under these conditions is called the best response to the move of your friend.

Let us start with two players. Ann’s **best response** to Beth’s move M is the move that yields the highest payoff for Ann, given Beth’s move M . There may be several best responses to a given move. To find Ann’s best response to Beth’s move M , we don’t even have to know Beth’s payoffs.

You find the best responses for the first player’s (Ann’s) moves by looking at the rows of the bimatrix one by one and selecting in each row the cell where the second entry is a maximum. The label of the corresponding column is the best response to the move corresponding to that row. In the same way, to find best responses against the second player’s (Beth’s) moves we consider the columns and pick in each column the cell with maximum first entry. The label of the corresponding row is the corresponding best response for the move corresponding to that column.

In the ADVERTISING example, the best response to advertising is to advertise, and the best response to not advertising is also to advertise. This holds for both players, since the game is symmetric.

In the TWO BARS example, the best response to a price of “2” is a price of “5”, the best response to a price of “4” is a price of “4”, and the best response to a price of “5” is a price of “4”. The game is symmetric.

Example 5 Let us give an asymmetric example. Assume Ann has four moves, A_1, A_2, A_3, A_4 , and Beth has three B_1, B_2 , and B_3 . The payoff bimatrix of this non zero-sum two-person game is

	B_1	B_2	B_3
A_1	1, <u>3</u>	2, 2	1, 2
A_2	<u>2</u> , <u>3</u>	2, <u>3</u>	2, 1
A_3	1, 1	1, <u>2</u>	<u>3</u> , <u>2</u>
A_4	1, 2	<u>3</u> , 1	2, <u>3</u>

We find Beth's best response to Ann's move A_1 by finding the largest second value (Beth's payoff) in the first row, which is underlined. That implies that Beth's best response to Ann's move A_1 is move B_1 . In the same way we underline the highest second values in other rows, and conclude that Beth's best responses to Ann's move A_2 are both moves B_1 and B_2 , Beth's best responses to move A_3 are both moves B_2 and B_3 , and that Beth's best response to move A_4 is move B_3 .

To find Ann's best responses, we underline in each column the highest first (Ann's payoff) entry. Therefore Ann's best response to Beth's move B_1 is A_2 , Ann's best response to B_2 is A_4 , and Ann's best response to B_3 is A_3 .

Best Response for Three Players

Best responses make also sense for games with three or more players. For detecting the best response moves of Beth, we look at the second entries (Beth's payoffs) in each column and mark the highest value. To detect the best response moves for the third player (let's call her Cindy) we look at the third entries of the rows and mark in each row the highest entry. For Ann the method is a little more complicated to explain. Here we look at first entries only, and compare cells having the same position in the different matrices, as "upper left", for instance.

Example 6 Let's find best responses in an example of a simultaneous three-person game where each player has two options, Ann has the moves A_1 and A_2 , Beth has B_1 and B_2 , and Cindy has C_1 and C_2 . Assume the payoffs are

		A_1				A_2	
		C_1	C_2			C_1	C_2
B_1	0, <u>2.1</u> , 0	-1, <u>1.1</u> , <u>0.1</u>		B_1	<u>0.1</u> , <u>1.1</u> , <u>1</u>	<u>1.1</u> , 0.1, -0.9	
B_2	<u>1</u> , 0, -1	0, 1, <u>1.1</u>		B_2	-0.9, 1, 0	<u>0.1</u> , <u>2</u> , <u>0.1</u>	

Because the highest second entry in the first column is 2.1, it is underlined. The highest second entry in the second column is 1.1, in the third column (first column of the second matrix) 1.1, and in the fourth column 2, so they are underlined. For Cindy's best responses, the highest third entry in the first row of the first matrix is 0.1. The highest third entry in the second row of the first matrix is 1.1. For the second matrix, the highest third entry in the first row is 1, and in the second row 0.1. For Ann, the highest first entry of upper-left cells in the two matrices is 0.1, the highest first entry of upper-right cells is 1.1, and we get 1 respectively 0.1 for the lower-left respectively lower-right cells.

2.2.4 Nash Equilibria

In this section we will identify outcomes—combinations of moves for each player—that are more likely to occur than others. An outcome is called a pure Nash equilibrium provided nobody can gain a higher payoff by deviating from the move, when all other players stick to their choices. A higher payoff for a player may be possible, but only if two or more players change their moves. An outcome, a combination of moves, is a **pure**

Nash equilibrium if each move involved is the best response to the other moves. A cell in the normal form is a pure Nash equilibrium if each entry is marked (underlined in our examples) as being the best response to the other moves. Nash equilibria were introduced by John Nash around 1950.

Nash equilibria are self-enforcing agreements. If some (non-binding) negotiation has taken place before the game is played, each player does best (assuming that the other players stick to the agreement) to play the negotiated move.

In the first half of the book, Nash equilibria will be pure. Chapter 27 will introduce mixed Nash equilibria.

In Example 5, there are two Nash equilibria: (A_2, B_1) and (A_3, B_3) . In the symmetric TWO BARS example $(4, 4)$ is the unique pure Nash equilibrium. As another example we consider the famous PRISONER'S DILEMMA.

PRISONER'S DILEMMA Adam and Bob have robbed a bank and been arrested. They are interrogated separately. Adam and Bob have the option to confess (move C) or to remain silent (move S). The police have little evidence, and if both remain silent they will be sentenced to one year on a minor charge. Therefore the police interrogators propose a deal: if one confesses while the other remains silent, the one confessing goes free while the other is sentenced to three years. However, if both talk, both will still be sentenced to two years. If each player's payoff is 3 minus the number of years served in jail, we get the following payoff bimatrix:

	S	C
S	2, 2	0, 3
C	3, 0	1, 1

It seems obvious that both should remain silent, but that's not likely to happen. Each player's move C strictly dominates move S . Furthermore, the best response to move S is C , and the best response to move C is also move C , therefore the pair (C, C) —both confessing forms the unique Nash equilibrium of this game.

The choice C —confessing—with payoffs of only 1 may seem counterintuitive if negotiations can take place in advance, but their terms are non-binding and cannot be enforced. It would be useless to agree on move S in advance, since each of the players would feel a strong urge to deviate (cheat). Only if binding agreements are possible, would both agree on the S - S combination, reaching a higher payoff. Thus PRISONER'S DILEMMA gives a paradoxical result. Players will play moves that result in lower payoffs for both than are possible. This is in part because the rules of the game do not allow binding agreements.

Not every simultaneous game has a (pure) Nash equilibrium. An example is Example 2 ROCK-SCISSORS-PAPER.

Next we consider a game with more than one Nash equilibrium:

Example 7 BATTLE OF THE SEXES: A couple, Adam and Beth, decide independently whether to go to a soccer game or to the ballet in the evening. Each person likes to do something together with the other, but the man prefers soccer, and the woman prefers ballet.

To simplify the game, we assume that the total payoff for each player is the sum of the payoffs (in terms of satisfaction) of being at the preferred place, which gives a satisfaction of c satisfaction units, and

being together with the partner, giving d satisfaction units. We have two variants, depending on whether c or d is larger, the low or high love variants. The payoff here is satisfaction instead of money. The assumption of the additivity of satisfaction is severe—satisfaction could just as well be multiplicative, or some more complicated function of c and d . It could even be that satisfaction in one area could interfere with satisfaction in the other. The satisfactions may differ for both persons, one appreciating the presence of the other more than the other, or one having a clear preference for soccer or ballet, when the other is indifferent. Examples will be given in the exercises.

As this example was devised before there were cell phones, we assume that no previous communication is possible. Here are the payoff bimatrices for both variants, where Adam chooses the rows and Beth chooses the columns.

High Love version, $c = 1, d = 2$			Low Love version, $c = 2, d = 1$		
	soccer	ballet		soccer	ballet
soccer	3, 2	1, 1	soccer	3, 1	2, 2
ballet	0, 0	2, 3	ballet	0, 0	1, 3

The high love version of BATTLE OF THE SEXES has two Nash equilibria: (soccer, soccer) and (ballet, ballet). For Adam choosing “soccer”, Beth’s best response is “soccer”. For Adam choosing “ballet”, Beth’s best response is “ballet”. Also, Adam’s best response for Beth choosing “soccer” is “soccer”, and his best response for Beth choosing “ballet” is “ballet”. The low love version has one Nash equilibrium, namely (soccer, ballet): both players go where they want to go anyway.

Modeling Note We made a simple assumption in the example, namely that the total payoff for a player is the sum of the utilities of certain ingredients. In many situations we will use this approach, since it is simple and is the way money is added. However, there are situations where additivity is not appropriate. One asset may establish worth only when combined with another asset, as a left shoe and the corresponding right shoe, or money and free time available. In many situations each has real value only in combination with the other.

Games with more than one pure Nash equilibrium are sometimes called “coordination games”, since if pre-game negotiations are allowed, the players have to agree on one of them. The high love version of BATTLE OF THE SEXES is an example. In this case, the obvious question is: which Nash equilibrium is the best? One idea is to concentrate on Pareto-optimal Nash equilibria only. A Nash equilibrium is **Pareto-dominated** by another Nash equilibrium if every player’s payoff in the first one is smaller or the same as in the second one. Nobody would object to move to the second Nash equilibrium. A Nash equilibrium is **Pareto-optimal** if it is not Pareto-dominated by any other Nash equilibrium, except maybe by some having exactly the same payoffs. In the BATTLE OF THE SEXES example, both Nash equilibria are Pareto-optimal.

For games with more than two players, we use the marking (underlining) procedure as described in the section on best responses. Then the cells with all entries underlined are the pure Nash equilibria.

In Example 6, a 3-player game, we have two pure Nash equilibria—the cells where all entries are underlined, where each move is the best response to the pair of moves of the other two players. These are the triples (A_2, B_1, C_1) and (A_2, B_2, C_2) . So player A will probably choose A_2 .

In our other example of a 3-player game, LEGISLATORS’ VOTE, let us underline the best responses:

A votes for a raise			A votes against a raise		
	C votes for raise	C votes against it		C votes for raise	C votes against it
B votes for raise	1, 1, 1	<u>1</u> , <u>1</u> , <u>2</u>	B votes for raise	<u>2</u> , <u>1</u> , <u>1</u>	0, -1, 0
B votes against	<u>1</u> , <u>2</u> , <u>1</u>	-1, 0, 0	B votes against	0, 0, -1	<u>0</u> , <u>0</u> , <u>0</u>

Here we have four pure Nash equilibria: the three outcomes where two legislators vote for a raise and one votes against, and the one where all three vote against. The fourth equilibrium is Pareto-dominated by the other three, so it is not Pareto-optimal and is therefore less important than the other three.

The next example, a 5-player game, illustrates how you can determine whether an outcome is a Nash equilibrium when you don’t have a bimatrix representation:

Example 8 5 KNIGHTS: Five knights, A, B, C, D, E, are electing their leader. Each one has a list of preferences. Examples of preferences, given from highest to lowest, are

A: A, D, E, C, B

B: B, C, E, A, D

C: C, E, D, B, A

D: D, B, C, E, A

E: E, C, B, A, D.

They elect in rounds. In each round, each knight submits one name. A knight is elected if he gets more votes than all the others. So even two votes may suffice if no other knight gets two votes. If no one is elected, we proceed to the next round. There are two versions:

Early Case If the knight’s first choice is elected, this is a payoff of 2 for that knight. If his second choice is elected, his payoff is 1. If nobody is elected and we proceed to the next round, the payoff is 0. If his third, fourth, or fifth choice is elected, his payoff is -1, -2, or -3.

Exhausted Case The knight’s first, second, and third choice gives payoffs of 2, 1, and 0. If no one is elected and we proceed to the next round, the payoff is -1. If his fourth or fifth choice is elected, his payoff is -2 or -3.

Each preference pattern defines a new game.

Because every player has five options, there are $5 \cdot 5 \cdot 5 \cdot 5 \cdot 5 = 3125$ outcomes. We could represent them with payoffs on a 5-dimensional cube.

Let’s instead look at an outcome and determine whether it is a Nash equilibrium in the two versions of the game. Assume A votes for A, B votes for B, C votes for C, D votes for C, and E votes for C. Then C is elected, and the payoffs for A, B, C, D, E are -2, 1, 2, -1, 1 in the early case game. Knight A is not happy but still has no reason to vote differently—if he voted for A or D instead, C would still be elected. But this outcome is not a Nash equilibrium, since D, knowing the voting pattern of the others, would rather have voted for B to obtain a tie and a payoff of 0.

In the exhausted case game, the payoffs for A, B, C, D, E for the same voting pattern are $-2, 1, 2, 0, 1$. Knight D still doesn't prefer C, but is now just content that somebody has been elected. That outcome is a Nash equilibrium in this version of the game. Nobody would, given the voting of the others, reconsider and vote differently. Knights A and D are still not happy, but they cannot unilaterally change this.

Let me show how we could search for Nash equilibria in the “exhausted knight” version. The idea is to start with any outcome, defined by a set of choices of the players. If all players are playing a best response to the other players' moves, we have a Nash equilibrium. Otherwise, at least one player does not play a best response yet—we let this player reconsider and play a best response. Then we evaluate the outcome again. Either we have a Nash equilibrium now, or we still have a player not playing the best response to the other players' moves. We continue, until we get a Nash equilibrium.

Look at the outcome where everybody votes for himself first. This would give a tie and everyone would prefer if his second choice would be elected. So, let's say D reconsiders and votes for B instead of for himself. Then B would be elected. B and E have better responses; A could vote for E instead for himself to get a tie and avoid the election of B. Now B, C, and E have better responses. Let's assume B plays his best response E to the other's moves. This voting pattern EECBE turns out to be a Nash equilibrium.

The process can be simulated in the [ExhaustedKnights](#) applet. Initially everybody votes his first preference. Change D's vote to B, then A's vote to E, then B's vote to E.

The process does not always terminate; confirm the following in the applet. We start with everyone voting for himself. Then A chooses a best response and votes for D. Then B chooses a best response and votes for C. After that D reconsiders and votes for B, then B reconsiders again, voting for himself again, and D reconsiders again, voting for himself again. After this we have an outcome that we had discussed earlier (voting pattern DBCDE) and the process could continue the same way forever.

Historical Remark In the Ph. D. thesis he wrote at Princeton University in 1950 the mathematician John Forbes Nash Jr. defined the equilibrium concept which is named after him. Later he did extraordinary work in other areas of mathematics. Around 1959 he became ill, suffering from paranoid schizophrenia throughout the 60s and the 70s. Surprisingly he recovered in the 80s. There is no Nobel prize in mathematics, but in 1994, with Reinhard Selten and John Harsanyi, Nash was awarded the Nobel prize in economics (to be more precise, the Nobel Memorial Prize in Economic Sciences). The award was given for his early work in game theory, including his definition of Nash equilibria and the existence theorem for them. The story of his life has been told in the book *A Beautiful Mind* by Sylvia Nasar [N1998], which in 2002 was made into an Oscar-winning movie with the same title.

2.3 Additional Topics

2.3.1 Best Response Digraphs

For a 2-player game, the best response information can be displayed in a graph. The bipartite **best response digraph** for two-player games is defined as follows: for every move of Ann we draw a white circle and for every move of Beth we draw a black circle. The circles are called the **vertices** of the digraph. From every white vertex we draw an arrow, an **arc**, towards black vertices that are best responses to the corresponding move of Ann. In the same way, arcs are drawn from black vertices towards best response white vertices. For Example 5, the best response digraph is shown in Figure 2.1.




Figure 2.1. The best response digraph for Example 5

Condensed Best Response Digraphs for Symmetric Games

In symmetric games, like ADVERTISING and TWO BARS, it suffices to display a condensed version of the best response digraph. For every move (of either player—the game is symmetric, therefore both players have the same moves as options) we draw one vertex, and we draw an arc from move X to move Y if Beth's Y is a best response to Ann's X (and therefore also Ann's move Y is a best response to Beth's move X). See Figures 2.2 and 2.3 for the best response digraph and the condensed best response digraph for the TWO BARS example. We may have curved arcs in the condensed version (in our example from vertex 4 to itself) when one of Ann's move (in our example move 4) is the best response to the corresponding move (move 4) of Beth.



Figure 2.3. The condensed best response digraph for the symmetric TWO BARS game

For two-player games, Nash equilibria can be recognized from the best response digraph. Any pair of moves with arcs between them, one being the best response of the other, is a Nash equilibrium. For symmetric 2-person games, in the condensed best response digraph any pair of arcs between two vertices, or any loop starting at a vertex and pointing to itself represents a pure Nash equilibrium. Those stemming from loops are symmetric insofar as both players use the same move in them. In symmetric games they may seem more natural than asymmetric Nash equilibria.

2.3.2 2-Player Zero-sum Symmetric Games

For the class of games with these three attributes, Nash equilibria, if they exist, can be spotted easily, and the maximin point of view is the same as the Nash equilibrium view.

Theorem *In every symmetric zero-sum simultaneous game,*

1. *every pure Nash equilibrium has zero payoff for both players, and*
2. *every maximin move of Ann with security level 0 versus any maximin move of Beth with security level 0 forms a Nash equilibrium.*

Our first theorem! What is a theorem anyway? So far this chapter has contained mostly definitions, examples, and facts about examples, such as the fact that TWO BARS has one Nash equilibrium. **Theorems** are also facts, not about single concrete examples but about general abstract mathematical objects like simultaneous games.

We want to provide a **proof** for this theorem. Although proofs can be complicated, they just provide the reasons why the theorem is true. You can accept the truth of a theorem based on the authority of the author or teacher, so it is all right if you skip the (very few) proofs in this book on first reading. But if you want to understand a mathematical area, you also have to understand the reasoning behind the proofs, at least to some extent.

Proof

1. Look at an outcome where one player, say Ann, has a payoff of less than zero. Then the move chosen by Ann could not be her best response for the move chosen by Beth, since she can always get 0 by choosing the same move as Beth.
2. Ann's minimum payoff in each row cannot exceed 0, since, if both players choose the same option, both have a payoff of 0. Therefore the security level cannot be larger than 0.

If Ann's minimum payoff in a row is less than 0, then each of Beth's best responses to the move of Ann corresponding to the row carries a payoff of more than 0 for Beth, therefore this move of Ann cannot be part of a Nash equilibrium by (1).

Therefore, if the security level is less than 0, there are no pure Nash equilibria.

If the security level (for both players, it is a symmetric game) equals 0, look at any maximin move for Ann and any maximin move for Beth. Then Ann's payoff in this move combination is at least 0, and Beth's payoff is at least 0. Since the game is zero-sum, both payoffs must be equal to 0. Then each move is the best response to the other move, and the move pair forms a pure Nash equilibrium.

It follows that a 2-player zero-sum symmetric game has no pure Nash equilibria provided the security level is less than 0. An example is ROCK-SCISSORS-PAPER.

One feature used in the analysis of simultaneous games is still missing. It is the topic of mixing moves and is discussed in Chapter 27.

Exercises

1. a) Write the matrices of the SIMULTANEOUS LEGISLATORS VOTE game in the variant where each of the three voters has also the option to abstain. The raise passes only if more agree than voting against. The loss of face by abstaining is relatively small, only \$200.
b) Solve the game, using the approaches discussed above.
2. Consider the following two-player game.

	L	M	R
U	1, 1	3, 4	2, 1
M	2, 4	2, 5	8, 1
D	3, 3	0, 4	0, 9

- Find the maximin moves for both players.
- Which moves are dominated?
- Find the bimatrix obtained by IESD.
- Find the bimatrix obtained by IEWD.

- Mark all best responses.
- Are there any Nash equilibria?

3. Analyze the following two-person zero-sum games for maximin moves, domination, best responses, and Nash equilibria:

a)

	L	R
U	1	2
D	3	4

b)

	L	R
U	1	2
D	4	3

c)

	L	R
U	1	3
D	2	4

d)

	L	R
U	1	3
D	4	2

e)

	L	R
U	1	4
D	2	3

f)

	L	R
U	1	4
D	3	2

4. Consider a two-person variant of the GUESS THE AVERAGE game: Ann and Beth simultaneously submit a number, 1, 2, 3, or 4. The player whose number is closest to $2/3$ of the average of both numbers gets \$1. Create the payoff bimatrix. Decide whether the game has a Nash equilibrium.
5. In the TWO BARS example a lack of tourists increases competition. Assume the number of natives is 4000. For which number of tourists would both bars choose \$4 as the price for a beer? For which tourist numbers is \$2 possible, and for which tourist numbers is \$5 possible?

6. Write the payoff bimatrix of the following game. Find maximin moves, domination, best responses, and pure Nash equilibria.

SCHEDULING A DINNER PARTY: Ann and Beth are not on speaking terms, but have a lot of common friends. Both want to invite them to a dinner party this weekend, either Friday or Saturday evening. Both slightly prefer Saturday. If both set the party at the same time, this will be considered a disaster with a payoff of -10 for both. If one plans the party on Friday and the other on Saturday, the one having the Saturday party gets a payoff of 5, and the other of 4.

7. Analyze the following game. Create payoff bimatrices consistent with the information given. Explain your choices. Then find the maximin moves, domination, and all pure Nash equilibria.

SELECTING CLASS: Adam, Bill, and Cindy are registering for a foreign language class independently and simultaneously. The available classes are ITA100 and FRE100. They do not care much which, but they care with whom they share the class. Bill and Cindy want to be in the same class, but want to avoid Adam. Adam wants to be in the same class as Bill or Cindy, or even better, both.

8. DEADLOCK: Two players play a symmetric game where each can either cooperate or defect. If they cooperate, both get an payoff of 1. If they defect, both get a payoff of 2. If one cooperates but the other defects, the one cooperating gets a payoff of 0, and the one defecting a payoff of 3.

Draw the bimatrix of the game. Find the maximin moves, possible domination, best responses, and find all pure Nash equilibria.

9. STAG HUNT: Two players play a symmetric game where each can hunt either stag or hare. If both hunt stag, both get an payoff of 3. If both hunt hare, both get a payoff of 1. If one hunts stag and the other hare, the stag hunter gets a payoff of 0, and the hare hunter a payoff of 2.

Draw the bimatrix of the game. Find the maximin moves, possible domination, best responses, and find all pure Nash equilibria.

10. CHICKEN: Two players play a symmetric game where each can either play dove or hawk. If both play dove, both get an payoff of 2. If both play hawk, both get a payoff of 0. If one plays dove and the other hawk, the one playing dove gets a payoff of 1, and the other one a payoff of 3.

Draw the bimatrix of the game. Find the maximin moves, possible domination, best responses, and find all pure Nash equilibria.

11. BULLY: Two players play the following game

	cooperate	defect
cooperate	2, 1	1, 3
defect	3, 0	0, 2

(compare [P1993]).

Find the maximin moves, possible domination, best responses, and find all pure Nash equilibria.

12. Two cars are meeting at an intersection and want to proceed as indicated by the arrows in Figure 2.4. Each player can proceed or move. If both proceed, there is an accident. A would have a payoff of -100 in this case, and B a payoff of -1000 (since B would be made responsible for the accident, since A has the right of way). If one yields and the other proceeds, the one yielding has a payoff of -5 , and the other one of 5 . If both yield, it takes a little longer until they can proceed, so both have a payoff of -10 . Analyze this simultaneous game, draw the payoff bimatrix, and find pure Nash equilibria.




Figure 2.4. Two cars at a crossing

13. Three cars are meeting at an intersection and want to proceed as indicated by the arrows in Figure 2.5. Each player can proceed or move. If two with intersecting paths proceed, there is an accident. The one having the right of way has a payoff of -100 in this case, the other one a payoff of -1000 . If a car proceeds without causing an accident, the payoff for that car is 5 . If a car yields and all the others intersecting its path proceed, the yielding car has a payoff of -5 . If a car yields and a conflicting path car as well, it takes a little longer until they can proceed, so both have a payoff of -10 . Analyze this simultaneous game, draw the payoff bimatrices, and find all pure Nash equilibria.




Figure 2.5. Three cars at a crossing

14. Solve the SIMULTANEOUS ULTIMATUM GAME. Display the payoff bimatrix, and investigate maximin moves, domination, best responses, and whether there are any equilibria.

15. Analyze a version of the BATTLE OF THE SEXES example where one partner has high love and the other low love. For the high love partner, being together with the partner is more important than being at the preferred location, whereas for the low love partner it is the opposite. Are there Nash equilibria?

16. Assume that a simultaneous two-player game has the best response digraph shown in Figure 2.6.




Figure 2.6. A best response digraph

- Display a possible payoff bimatrix. Can you find a zero-sum payoff bimatrix generating this best response digraph?
17. In the 5 KNIGHTS game described in Example 8 with preferences as described there, determine whether the voting pattern ECEDE (A votes for E, B votes for C, etc.) forms a Nash equilibrium in the early case game or in the exhausted case game.
18. In the 5 KNIGHTS game described in Example 8 with preferences as described there, determine whether the voting pattern ACEBE (A votes for A, B votes for C, etc.) forms a Nash equilibrium in the early case game or in the exhausted case game.
19. Use the 5KNIGHTS applet to find a Nash equilibrium for the early version of the 5 KNIGHTS game.

Project 1

Reacting fast or slow Assume the five players of the 5 KNIGHTS game are negotiating what to vote before actually voting. They start with the obvious proposal of everybody voting for his first choice. If somebody's move is not the best response to the other players' moves he changes the proposal, proposing his best response as his move. This is repeated until a Nash equilibrium is found.

The process is not unique, since if two or more players want to reconsider, only one will modify the proposal at a time. Discuss whether it is better in such a negotiation to always modify early, or to wait and see whether the others change the proposal first. Simulate the process for different preferences in the 5KNIGHTSRANDOM applet with the assumption that A always reacts faster than B, B always reacts faster than C, and so on. Do this at least 30 times, and keep track how often each one of the players is elected in the resulting Nash equilibrium (which you hopefully get—there may also be cyclic cases as discussed in Example 8 of Chapter 27).

CHAPTER 3

Example: Selecting a Class

Prerequisites: Chapters 1 and 2.

All students in Soap College have to enroll either in FRE100 or ITA100 in their second semester. We assume that every student prefers one of them. There is a group of students, however, the “drama queens and kings”, who have strong likes and dislikes between their members, and for whom it is most important to be in a class with as many as possible of drama queens and kings they like and with as few as possible they dislike. A measure of this is the difference between the number of liked and the number of disliked drama queens or kings in the class. Only if this difference is the same for both classes would the drama queens or kings make their decisions according to their personal preferences for a language.

The dean at Soap College is concerned that many of the group members will end up in a class for social reasons, and assigns to us the task of investigating this problem.

The drama queens and kings don’t care about other students in the course. They have strong feelings only about their fellow drama queens and kings. Group members could also be indifferent towards other group members. We can model this as a game by assuming that each group member gets a payoff of 9 for every group member he or she likes who is in the same class, and a payoff of -9 for every group member he or she dislikes who is in the same class, and an additional payoff of 1 if he or she gets the course for which he or she has a preference. The total payoff is the sum of the numbers.

We will discuss several examples in this chapter, mostly where the group of drama queens and kings consists of three to six persons, and for two or three available classes. For three players we can create and use the payoff matrices as described in Chapter 2 to find Nash equilibria. We will see that for more than three players the matrices may be too complicated, but we can find Nash equilibria by best response dynamics. This is the only place where a game for more than three players is discussed.

3.1 Three Players, Two Classes

There are many variants of the game, depending on liking, disliking, and indifference among the three players, and also their preferences for Italian or French. Given a pattern of liking and disliking, there are eight possibilities, and so eight different games, of class preferences of the three players. In this section we will discuss all eight.

3.1.1 “I like you both”

The three group members are Adam, Beth, and Carl, and there is mutual sympathy between Adam and Beth, and between Beth and Carl. There is no disliking (yet) in this group.

Adam and Beth prefer French and Carl prefers Italian

Student Activity Three students should role-play Adam, Beth, and Carl, where Adam and Beth prefer French and Carl prefers Italian. Reserve one corner of the classroom for FRE100 and another one for ITA100. Students should tentatively move to the area they prefer. When everybody has chosen, each student should think about the two options, given the present location of the other two players and should move if this gives improvement.

The payoff matrices, with best responses underlined, are

A chooses FRE100			A chooses ITA100		
	C chooses FRE100	C chooses ITA100		C chooses FRE100	C chooses ITA100
B chooses FRE100	10, <u>19</u> , <u>9</u>	<u>10</u> , <u>10</u> , 1	B chooses FRE100	0, <u>10</u> , <u>9</u>	0, 1, 1
B chooses ITA100	1, 0, 0	1, 9, <u>10</u>	B chooses ITA100	<u>9</u> , 9, 0	<u>9</u> , <u>18</u> , <u>10</u>

The game has two Nash equilibria, with all players selecting the same class. If all choose French, then two have chosen the class they should go to (according to the dean who has only academics in mind), and if all choose Italian, one has made the right choice. On average we have 1.5 right choices, which is not better than random.

Other class preferences

It doesn't really matter which classes Adam, Beth, and Carl prefer. Depending on the preferences, the payoff matrices are

A chooses FRE100			A chooses ITA100		
	C chooses FRE100	C chooses ITA100		C chooses FRE100	C chooses ITA100
B chooses FRE100	<u>9+ε</u> , <u>18+ε</u> , <u>9+ε</u>	<u>9+ε</u> , 9+ε, 0+ε	B chooses FRE100	0+ε, 9+ε, <u>9+ε</u>	0+ε, 0+ε, 0+ε
B chooses ITA100	0+ε, 0+ε, 0+ε	0+ε, 9+ε, <u>9+ε</u>	B chooses ITA100	<u>9+ε</u> , 9+ε, 0+ε	<u>9+ε</u> , <u>18+ε</u> , <u>9+ε</u>

where ϵ is 0 or 1, depending on whether the student is in the right class. The bold and underlined numbers are best responses. The bold but not underlined numbers could be best responses, but don't have to be, depending on the class preferences of the players (thus depending on which ϵ s are 0 and which are 1). But whether these outcomes are best responses or not, we get the two Nash equilibria of all three students going to the same class.

For this analysis we do not really need the matrices. Just look at the class activity above. If Adam or Carl are tentatively not in the same class as Beth, they will move, not matter what class they would slightly prefer. Beth may also move if Adam and Carl are both in the other class, or if Adam and Carl are presently in different classes and Beth is not in the class she would slightly prefer. Since Beth eventually settles, Adam and Carl will move to the class where Beth finally is, and will settle there as well. All three resting means we have a Nash equilibrium.

Thus the disappointing outcome for the dean is: in the “I like you both” case, on average only 50% of the students will end in a class they would prefer if no personal feelings were involved.

3.1.2 Disliking the Rival

The three group members are Adam, Beth, and Carl, and there is mutual sympathy between Adam and Beth, and between Beth and Carl, and mutual antipathy between Adam and Carl.

Student Activity For this pattern, students should do the simulation where Adam and Beth prefer French and Carl prefers Italian.

Student Activity Now try the simulation where Adam and Carl prefer French and Beth prefers Italian.

Adam and Beth prefer French and Carl prefers Italian

A chooses FRE100			A chooses ITA100		
	C chooses FRE100	C chooses ITA100		C chooses FRE100	C chooses ITA100
B chooses FRE100	1, 19, 0	10, 10, 1	B chooses FRE100	0, 10, 9	-9, 1, -8
B chooses ITA100	-8, 0, -9	1, 9, 10	B chooses ITA100	9, 9, 0	0, 18, 1

Here everybody goes to the class he or she should (according to the dean) attend.

Adam and Carl prefer French and Beth prefers Italian

A chooses FRE100			A chooses ITA100		
	C chooses FRE100	C chooses ITA100		C chooses FRE100	C chooses ITA100
B chooses FRE100	1, 18, 1	10, 9, 0	B chooses FRE100	0, 9, 10	-9, 0, -9
B chooses ITA100	-8, 1, -8	1, 10, 9	B chooses ITA100	9, 10, 1	0, 19, 0

We get three Nash equilibria. All three cases where two of them go where they should go according to their class preferences, and the third one does not, are possible. If Beth stays in French, then it is because both Adam and Carl are there. If Adam chooses Italian, then the reason is that Beth takes it too whereas Carl does not. There is a similar outcome in the third case.

The other cases

If all three students prefer French, then all will go there. If Adam prefers French but Beth and Carl prefer Italian, then also everybody will go to the class he or she should go.

Overall, “Disliking the rival” would please the dean. In most cases, students go where they should go.

3.1.3 Outsider

Assume that Adam and Beth like each other, and that Carl likes both of them, whereas they both dislike Carl. This is an awkward situation for poor Carl, but also for the game, as we will see.

Student Activity For this pattern, students should do the simulation where Adam and Carl prefer French and Beth prefers Italian.

What happens is that we don’t have any Nash equilibrium. So students will never stop moving in this simulation.

3.2 Larger Cases

The drama queens and kings group could have more than three members, but there could also be more than two classes to choose from, as for instance ITA100, FRE100, and GER100. Then the number of positions increases from $8 = 2^3$ to $2^5 = 32$ for two classes and five persons, or $3^6 = 729$ for three classes and six persons. For these examples, checking all possible positions to see whether they form Nash equilibria is not a problem if you use a computer, but it is a lot of work by hand.

Nash equilibria can be found automatically during registration. In most colleges students can tentatively register for a class, and change their decisions later. While changes are possible, students can also see how other students have decided. Students who can profit from changing class will move to their best response option. We assume that they look at the class lists and make their changes one-by-one, in a random order. Only if there is a Nash equilibrium will no change occur. The naive hope may be that adjusting moves to best responses repeatedly leads to better and better positions, satisfying more and more players, until eventually a Nash equilibrium is found. This hope is not always justified. There are even cases where this process does not terminate, even if the game has a Nash equilibrium. However, practically this process often does find some Nash equilibrium.

Assume there are five persons, A, B, C, D, E, with liking from A to B, from B to C, from C to D, from D to E, and from E to A, and with disliking from B to E, from E to D, and from D to B, and with mutual disliking between A and C. Moreover B, C, and D have slight preferences for FRE100, and A and E slightly prefer ITA100. This is indicated in Figure 3.1, where full arrows indicate liking, dotted arrows indicate disliking, and the class preferences are attached to the student names.




Figure 3.1. Liking (arrows) and disliking (dotted arrows) between five persons with class preferences indicated

Every student likes one other student of the group, and dislikes another one. This implies that in a distribution of the students to the classes, a student X

- wants to change classes if the disliked one is in the same class and the liked one is not, but
- is satisfied with staying if the liked one is in the same class, but the disliked one is not.
- follows her or his slight class preference if both liked student and disliked student are in the same class.

Student Activity For this pattern, five students should do the simulation.

Student Activity Instead, you can simulate the process in the [Selecting52A](#) applet. The data of our example are already filled in. Don't click any buttons. Just click the faces until every face smiles.

Confirm that you get a Nash equilibrium if A, D, and E choose ITA100 and B and C choose FRE100. This is the only one. Did you find it in the simulation(s)?

The last question is not just rhetorical—you may have not. Go back to the applet and move A, E, B, A, E, B, A, E, B, and so on. Six positions repeat over and over, without resulting in the Nash equilibrium.

Student Activity Use the [Selecting52A](#) applet and the same method to check for Nash equilibria in the modification where the likings and dislikings between group members remain the same, but the class

preferences are different. You can create other class preferences either by clicking the “random class data” button, or by changing the numbers in the text fields and clicking the “read data” button.

3.3 Assumptions

How realistic are the assumptions? Do we always like everybody just the same, with a value of 9 in our models? Obviously not, but this can be made more realistic by assigning different liking payoffs, as we can for class preferences. More serious is the question of exactness of the values. Let’s say A likes B with value 9, and C with value 4. Can we be sure that the values are not 8 and 5? Or 8.9 and 4.1? With slightly changed values of the data, can we still be sure to have the same Nash equilibrium?

You can investigate these questions in the applets [Selecting52](#), [Selecting53](#), [Selecting62](#), and [Selecting63](#), where the payoffs for having a person in the same class, or being in a certain class, can be changed. Random data can also be generated.

Exercises

1. **Cyclic Liking:** Investigate two cases (based on their class preferences) of the situation where Adam likes Beth, Beth likes Carl, and Carl likes Adam, and otherwise there is indifference.
2. **Cyclic Disliking:** Analyze two class-preference cases of the situation where Adam dislikes Beth, Beth dislikes Carl, and Carl dislikes Adam, and otherwise there is indifference.
3. **Darling:** Now assume Adam and Carl both like Beth, but Beth does not reciprocate. Analyze two class-preference cases.
4. **Cyclic liking and disliking:** Let’s assume that Ann likes Beth but Beth dislikes Ann (such things happen), that Beth likes Cindy but Cindy dislikes Beth, and that Cindy likes Ann but Ann dislikes Cindy. Analyze two class-preference cases.
5. Are there any Nash equilibria if the group consists of five students with the liking pattern and class preferences in Figure 3.2?
6. Use the applet [Selecting52](#) to investigate the case of five persons and two classes, where C, D, and E like both A and B, and A and B dislike each other.
7. Use the applet [Selecting63](#) to investigate the case of six persons and three classes, where C, D, E, and F like both A and B, and A and B dislike each other.




Figure 3.2. Liking (arrows) and disliking (dotted arrows) between five persons with class preferences as indicated

Project 2

Assume there are five persons in the group, with liking and dislikings as in Figure 3.1, but with unknown class preferences. Can you tell what will happen, or what most likely will happen? Use the applet [Selecting52A](#) to find Nash equilibria (if they exist) for all 32 possible instances. Is there a pattern in the Nash equilibria?

Project 3

Are Nash equilibria more frequent in groups where sympathy occurs more often than antipathy than in groups with many conflicts or in groups with mixed sympathy and antipathy? Use applet [Selecting63](#) to create at least 30 examples for each case (where you still have to explain what more sympathy or more antipathy means, and how you created your examples) and calculate the Nash equilibria using the “Nash” button. What is the percentage of Nash equilibria, in each case? Explain your findings.

Project 4

For the same group of drama queens and kings, and the same chemistry within the group, are Nash equilibria more frequent if three classes are offered than if two classes are offered? Create at least 30 examples of sympathies and antipathies within a group of five persons. Then randomly assign class preferences for the two or three classes available or three classes available for each of the examples. Use the applets [Selecting53](#) and [Selecting52](#) to check whether Nash equilibria exist in both cases. You can also investigate the cases of six persons using the applets [Selecting63](#) and [Selecting62](#). Explain your findings.

CHAPTER 4

Example: Doctor Location Games

Prerequisites: Chapters 1 and 2.

In this chapter and its sibling on Restaurant Location Games we discuss the simultaneous versions of different location games played on undirected graphs.

Like a digraph, an undirected **graph** consists of vertices. But instead of being connected by arcs with directions, vertices are connected by undirected **edge** that can be used both ways. Usually they are drawn with curves connecting the vertices. Vertices connected by an edge are called **adjacent**. Edges may cross, as in the graph in Figure 4.1, where there is an edge connecting vertices 4 and 5, and one connecting vertices 3 and 6, but none connecting vertices 4 and 6. Therefore vertices 4 and 6 are not adjacent, nor are vertices 3 and 5.




Figure 4.1. A graph

We use graphs for modeling transportation networks. The vertices could represent towns, and the edges roads between the towns. Two towns are called **adjacent** if they are connected by a direct road. Crossings between edges have bridges, it is not possible to change the roads there.

4.1 Doctor Location

Doctor Location Game: On a small island, different towns, all having the same number of residents, are connected by a system of streets. Two medical doctors, Ann and Beth, are about to settle. On average, everyone on the island sees a doctor once a year, and always chooses the nearest one. The measure for distance is the number of roads one has to travel to go there. In case of a tie, half of the citizens of that town go to Dr. Ann and the other half to Dr. Beth.

In the simultaneous version of the game, both Ann and Beth are forced by the island government to submit a proposal for their choice of location simultaneously. In what towns should they choose to locate?

Try the game on different graphs in one of the applets [LocationDr2](#), [LocationDr4](#), [LocationDr5](#), [LocationDr6](#), or [LocationDr8](#).

4.1.1 An Example Graph

Consider as an example the graph displayed in Figure 4.2 and let me explain how the payoff matrix of the game is derived.




Figure 4.2. A graph

- If both players choose the same town, then they just split all islanders as potential patients, therefore both get a payoff of 3.5 towns each.
- Assume Ann chooses town 1 and Beth town 2. See Figure 4.3 for a visualization of what happens. Ann gets the whole town 1, and Beth the whole town 2. Town 3 is split, since both town 1 and 2 are adjacent. This is indicated by coloring town 3 half blue and half red. Towns 4 and 6 are also both split. From them one can reach both town 1 and town 2 in two steps, but cannot do it in one step. The whole town 5 is going to Dr. Ann, since town 1 is adjacent to town 5 but town 2 is not, and the whole town 7 is going to Dr. Beth. Counting all patients, in this outcome Ann and Beth get a payoff of 3.5 each.




Figure 4.3. Payoff for some placement on the graph

- Assume Ann chooses town 2 and Beth town 5. Again both Ann and Beth get a payoff of 3.5.
- If Ann chooses town 1 and Beth town 3, Ann gets a payoff of 4 and Beth gets 3.
- Assume Ann chooses town 6 and Beth town 3. Ann gets a payoff of 3 and Beth gets a payoff of 4.

You can check the computations in the applet [LocationDr6](#).

We don't need the whole bimatrix in order to check whether a pair of moves forms a Nash equilibrium or not. Usually we don't need the bimatrix at all. Consider a pair of moves where Ann gets more payoff than $n/2$ and Beth gets less, with n the number of vertices of the graph, as in the town 1 versus town 3 outcome in the above example. Why can't this form a Nash equilibrium? Beth would want to change her decision, and put her location at the same vertex as Ann, thereby increasing her payoff to $n/2$. This reasoning implies that all Nash equilibria would have to give a payoff of $n/2$ to both players.

So do all pairs of moves giving both players a payoff of $n/2$ form a Nash equilibrium? Not necessarily, not if one player can achieve more than $n/2$ by relocating. But since the security level cannot exceed $n/2$, this implies either that one of the moves is not a maximin move, or that some player’s security level is less than $n/2$.

Haven’t we done this reasoning recently? Remember the Theorem in subsection 2.3.2. Our doctor location games are not zero-sum games, but they are constant-sum games, since each islander goes to one doctor once a year. Thus the sum of the patients of both doctors equals the number of islanders. Constant-sum games behave like zero-sum games if we subtract half of the constant sum from each payoff. Therefore we get

Fact *In doctor location games with n vertices, the Nash equilibria are formed by every maximin move of Ann versus every maximin move of Beth, provided both security levels (guaranteed payoffs) are $n/2$.*

Even when doctor location games have several Nash equilibria and are coordination games, coordination is not necessary—it suffices that each player plays just one of their maximin moves.

Here is the payoff bimatrix for our example:

	1	2	3	4	5	6	7
1	3.5, 3.5	3.5, 3.5	4, 3	4.5, 2.5	3.5, 3.5	4, 3	4.5, 2.5
2	3.5, 3.5	3.5, 3.5	4, 3	4, 3	3.5, 3.5	4, 3	4.5, 2.5
3	3, 4	3, 4	3.5, 3.5	4, 3	3.5, 3.5	4, 3	4, 3
4	2.5, 4.5	3, 4	3, 4	3.5, 3.5	2.5, 4.5	3.5, 3.5	3.5, 3.5
5	3.5, 3.5	3.5, 3.5	3.5, 3.5	4.5, 2.5	3.5, 3.5	4.5, 2.5	4, 3
6	3, 4	3, 4	3, 4	3.5, 3.5	2.5, 4.5	3.5, 3.5	3.5, 3.5
7	2.5, 4.5	2.5, 4.5	3, 4	3.5, 3.5	3, 4	3.5, 3.5	3.5, 3.5

The maximin moves for both players are moves 1, 2, and 5, and the corresponding values, the security levels, are 3.5 for each. The four other moves 3, 4, 6, and 7 are weakly dominated by them. There are nine pure Nash equilibria, moves 1, 2, 5 versus moves 1, 2, 5.

4.1.2 No (Pure) Nash Equilibrium?

The graph in Figure 4.4 has no (pure) Nash equilibrium—both security levels are below 4.5. See [LocationDr30](#).




Figure 4.4. A graph without pure Nash equilibrium for the doctor location game

Here is the payoff bimatrix:

	1	2	3	4	5	6	7	8	9
1	4.5, 4.5	3, 6	4, 5	4.5, 4.5	3, 6	3.5, 5.5	4, 5	3.5, 5.5	4.5, 4.5
2	6, 3	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4, 5	5, 4	4.5, 4.5	4.5, 4.5
3	5, 4	4.5, 4.5	4.5, 4.5	6, 3	4, 5	4.5, 4.5	4.5, 4.5	5, 4	5.5
4	4.5, 4.5	4, 5	3, 6	4.5, 4.5	3.5, 5.5	3, 6	3.5, 5.5	4, 5	4.5, 4.5
5	6, 3	4.5, 4.5	5, 4	5.5	4.5, 4.5	4.5, 4.5	4.5, 4.5	4, 5	5, 4
6	5.5	5, 4	4.5, 4.5	6, 3	4.5, 4.5	4.5, 4.5	4, 5	4.5, 4.5	5, 4
7	5, 4	4, 5	4.5, 4.5	5.5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5	6, 3
8	5.5	4.5, 4.5	4, 5	5, 4	5, 4	4.5, 4.5	4.5, 4.5	4.5, 4.5	6, 3
9	4.5, 4.5	3.5, 5.5	3.5, 5.5	4.5, 4.5	4, 5	4, 5	3, 6	3, 6	4.5, 4.5

4.1.3 How Good are the Nash Equilibria for the Public?

The inhabitants of the island are not players of the game, since they don’t make decisions. Still there is a payoff for them, having a doctor close to their home or not. The sum of the lengths of paths from vertices to the nearest doctor could be seen as a measure of how good the location of the doctors is to society—the smaller the number, the better. Let’s call this number the distance sum, given the doctors’ locations. Finding a placement of two doctors minimizing the distance sum is tedious but not too difficult—all we have to do is to compute the distance sum for every combination. In the graph in Figure 4.3, if both doctors go to vertex 1, the distances from vertices 1, 2, 3, . . . and so on to the closest doctor (at vertex 1) are 0, 1, 1, 2, 1, 2, and 2. Therefore the distance sum for this placement is $0 + 1 + 1 + 2 + 1 + 2 + 2 = 9$. If one doctor is located at vertex 1 and the other at vertex 2, the distance sum is $0 + 0 + 1 + 2 + 1 + 2 + 1 = 7$. If the doctors are located at vertices 2 and 5, the distance sum is $1 + 0 + 1 + 1 + 0 + 1 + 1 = 5$, which is lowest possible. Letting the two doctors choose their location freely does not necessarily minimize this distance sum, and therefore does not maximize society’s payoff.

4.2 Trees

Now we know that we only need to look for vertices that guarantee half the number of vertices as payoff to find all Nash equilibria. Is it possible to identify them without performing calculations but by looking at the structure of the graph? Do we know how many vertices of this type we have? Do we know which graphs have pure Nash equilibria and which one do not?

In this section, we will answer these questions for special graphs, namely trees. A **tree** is a graph where between any pair of vertices there is exactly only path. Examples are the graphs in Figures 4.5 and 4.6.




Figure 4.6. Another tree

You can generate trees and play the doctor location game in the applet [LocationDrTree](#).

Why trees? The reason is that the maximin move vertices—those are all we have to find— can also be described by the branch-weight, as will be shown:

1. We need the following definitions: If we remove a vertex x from a tree, the tree falls into parts. These parts are called the “branches” at vertex x , or “ x -branches”, and the maximum number of vertices in an x -branch is called the **branch-weight** $bw(x)$ of x . In our examples, the branch-weights of the vertices are shown in Figures 4.7 and 4.8.




Figure 4.8. Another tree with branch-weights

2. The first observation is that the best response location to a location x is always either the same vertex or one of its neighbors. Because for every other placement of Beth she could improve by moving closer to Ann’s location x . Therefore from now on we consider cases of placements on the same or adjacent vertices.
3. If Ann is located at vertex x and Beth at vertex y , then all but one of the x -branches do not contain y . The vertices in the x -branches are customers of x . In the same way, all but one y -branches do not contain x , and the vertices in the y -branches are customers of y . Only the vertices between x and y are split between x and y . If x and y are not adjacent, then y could improve the number of customers by moving closer to x . The vertices in the former y -branches not containing x still belong to y , but there are some gains in the former “between” part just about in the middle between x and y . Consequently, the best response location to a location x is always either the same vertex or one of its neighbors. Therefore from now on we consider cases of placements on the same or adjacent vertices.
4. The case where both players choose the same vertex is not too interesting, since the payoffs are split equally. If Ann and Beth locate on adjacent vertices x and y , then removing the edge between them splits the tree into two parts. One is the y -branch attached to y at x , and the other the x -branch attached to x at y . This is the partition of the vertices into those visiting Ann’s office versus those visiting Beth’s office.
5. Assume that Ann places her office at vertex x in a tree with n vertices. Beth can always achieve a payoff of $n/2$ by placing at x too, but whenever Beth places at a neighbor of x , she gets just the corresponding branch as customers. So Beth can achieve the branch-weight $bw(x)$ of x as payoff by placing on a neighbor of x , but not more. Since we have a zero-sum game, Ann’s guaranteed payoff when placing at x is either $n/2$ or $n - bw(x)$, whichever is smaller. That implies that, provided there are some vertices of branch-weight less than or equal to $n/2$, all the vertices are the maximin moves for the players.

Next we will show that vertices of branch-weight less than or equal to $n/2$ exist, and that there are few of them [Z1968]:

1. If an x -branch of weight k is attached to x at vertex y , then $bw(y) \geq n - k$, since the y -branch attached to y at x contains all vertices outside of the latter branch.
2. Therefore all neighbors of a vertex x of branch-weight less than $n/2$ have branch-weight greater than $n/2$. Moreover, since a vertex of branch-weight $n/2$ can contain only one branch of this size, it has exactly one neighbor of branch-weight $n/2$, and all other neighbors have larger branch-weight.
3. If a vertex x has branch-weight greater than $n/2$, then the neighbor y of x in the large branch has smaller branch-weight, but all other neighbors z of x have larger branch-weights than x . For, the y -branch attached at x consists of x and all small x -branches, all vertices outside the large x -branch, and has therefore less than $n/2$ vertices. All other y -branches are strictly smaller than the large x -branch. For the second statement, the z -branch attached at x contains the large x -branch of x . Therefore, starting at any vertex, we can always move to neighbors with smaller branch weight until we arrive at a vertex with branch-weight $\leq n/2$.

Consequently we have:

Theorem *In a tree, the maximin move vertices are the vertices minimizing the branch-weight. They are either one vertex of branch-weight less than $n/2$, or two adjacent vertices of branch-weight $n/2$. The Nash equilibria are all pairs of these vertices.*

Therefore in the first tree example, the only maximin move is to play vertex 6. In the second tree example, there are two maximin moves, vertices 6 and 8, both with a branch-weight of 6.

If we had only one doctor to place, the locations minimizing the distance sum are called the **medians** and were first investigated by Camille Jordan for trees [J1869]. The distance sums for the vertices in our two trees are shown in Figures 4.9 and 4.10.




Figure 4.10. Another tree with distance sums to the vertices

Thus the medians are vertex 6 in the first tree, and vertices 6 and 8 in the second. The medians, the vertices minimizing the distance sum, are exactly the vertices for the maximin moves, the vertices with minimum branch-weight, in these two examples. Camille Jordan found that the median for trees is always either one vertex, or two adjacent vertices, and Bohdan Zelinka showed 199 years later that the median of a tree consists of exactly those (one or two) vertices minimizing the branch-weight [Z1968].

Therefore, the Nash equilibria locations for trees would be perfect for society if there were only one doctor, but for two doctors, they are almost always not optimal. Both doctors behave as if they were the only doctor there. This is a case where competition would not lead to a solution best for society. In the first example tree, for instance, both doctors will go to vertex 6, whose distance sum is 19. However, if one would go to

vertex 2 and the other to vertex 6, the distance sum would be 9! Surprisingly, in the second tree the Nash equilibrium of vertex 6 against vertex 8 is also optimal for society, since the distance sum equals 15, the same as for locations at vertices 6 and 8, locations at vertices 6 and 10, at vertices 4 and 8, or at vertices 4 and 10. This is because the tree is small. For larger trees, the Nash equilibrium location is usually not optimal for society, as it does not minimize the distance sum.

4.3 More than one Office (optional)

We change the game slightly. Each doctor maintains two offices, in two towns, one for before noon, and one for afternoons. We assume that people can wait half a day, so all four offices are options. In case of a tie, there are very slight preferences of the population towards the doctors, split equally, so each doctor could expect half a patient. The choices of the offices are revealed simultaneously.

You can experiment with different graphs in all applets mentioned. An example on a selection and the corresponding payoffs for the example graph considered before in [LocationDr6](#) is given in [Figure 4.11](#).




Figure 4.11. Two offices for each doctor

We have a symmetric simultaneous 2-player zero-sum game. But the number of options for each player is now larger than before. If the town graph has n vertices, then each doctor has n options with both her offices placed in the same town. These options are obviously dominated. Then there are $n \cdot (n - 1)/2$ options with both offices in different towns. Together, each doctor has $n + n \cdot (n - 1)/2$ options. For our 7-vertex example Graph 6, this adds to 28 options for each doctor. In the 10-vertex Graph 4, we have 55 options for each doctor, so we have a 55×55 matrix.

But we also still have a symmetric 2-player zero-sum game. So we know that every pair of maximin moves of the two players yields a pure Nash equilibrium, provided they carry a guaranteed payoff of half the number of vertices. If the maximin moves have a lower guaranteed payoff, there is no pure Nash equilibrium.

Exercises

1. For the graph in applet [LocationDr5](#), part of the payoff bimatrix is given below. Use the applet to fill the gaps in the matrix, and find the Nash equilibria, if there are any.

	1	2	3	4	5	6	7
1	3.5, 3.5	2.5, 4.5	3, 4	3.5, 3.5	2.5, 4.5	4.5, 2.5	3.5, 3.5
2	4.5, 2.5	3.5, 3.5	3.5, 3.5	—, —	—, —	—, —	4.5, 2.5
3	4, 3	3.5, 3.5	—, —	4, 3	3.5, 3.5	6, 1	4, 3
4	3.5, 3.5	—, —	3, 4	3.5, 3.5	2.5, 4.5	4.5, 2.5	3.5, 3.5
5	4.5, 2.5	—, —	3.5, 3.5	4.5, 2.5	3.5, 3.5	5, 2	—, —
6	2.5, 4.5	—, —	1, 6	2.5, 4.5	2, 5	3.5, 3.5	2, 5
7	3.5, 3.5	2.5, 4.5	—, —	—, —	2.5, 4.5	5, 2	3.5, 3.5

2. For the graph given in applet [LocationDr2](#), part of the payoff bimatrix is given below. Use this applet to fill the gaps in the matrix, and find the Nash equilibria, if there are any.

	1	2	3	4	5	6	7	8	9
1	4.5, 4.5	—, —	—, —	—, —	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5
2	—, —	4.5, 4.5	5, 4	5, 4	5, 4	5, 4	6, 3	5, 4	5, 4
3	—, —	4, 5	—, —	5, 4	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5
4	4, 5	—, —	4, 5	4.5, 4.5	4, 5	4.5, 4.5	—, —	4.5, 4.5	4, 5
5	4.5, 4.5	4, 5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5
6	4.5, 4.5	4, 5	4.5, 4.5	4.5, 4.5	4.5, 4.5	4.5, 4.5	—, —	4.5, 4.5	4.5, 4.5
7	4, 5	3, 6	4, 5	4, 5	4, 5	—, —	4.5, 4.5	—, —	4, 5
8	4.5, 4.5	4, 5	4.5, 4.5	4.5, 4.5	4.5, 4.5	4.5, 4.5	5.5, 3.5	4.5, 4.5	4.5, 4.5
9	4.5, 4.5	4, 5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5

3. Create the payoff bimatrix and find all Nash equilibria for the left graph in Figure 4.12.
4. Create the payoff bimatrix and find all Nash equilibria for the right graph in Figure 4.12.




Figure 4.12. Two graphs

5. Find all best responses for the graph in Figure 4.4. Show that there is no pure Nash equilibrium.
6. Find the median and all pure Nash equilibria of the corresponding doctor location game for the left tree in Figure 4.13.




Figure 4.13. Two trees

7. Find the median and all pure Nash equilibria of the corresponding doctor location game for the right tree in Figure 4.13.
8. Vertices that are adjacent to just one vertex, i.e. that have just one neighbor, are called **end vertices**. Draw a tree with fourteen vertices, with exactly five end vertices. Calculate the branch-weights of all vertices. Find the median and all pure Nash equilibria of the corresponding doctor location game.
9. Draw a tree with twelve vertices, with exactly six end vertices. Calculate the branch-weights of all vertices. Find the median and all pure Nash equilibria of the corresponding doctor location game.
10. Can the symmetric constant-sum game in Table 4.1 be formulated as a doctor location game?

	1	2	3
1	1.5, 1.5	2.5, 0.5	0.5, 2.5
2	0.5, 2.5	1.5, 1.5	2.5, 0.5
3	2.5, 0.5	0.5, 1.5	1.5, 1.5

Table 4.1.

Projects

A **cycle** is a graph with a number of vertices $a_1, a_2, \dots, a_n$, and an edge between a_1 and a_2 , an edge between a_2 and a_3 , and so on, until the edge between a_{n-1} and a_n and an additional edge between a_n and a_1 .

Project 5

Doctor location on MOPs A **MOP** or **maximal outerplanar graph** is obtained from a cycle by adding diagonal edges that do not cross each other until no more edges can be drawn without crossing. An example is given in Figure 4.14. Can you describe how many Nash equilibria are possible for doctor location games on MOPs? Start by analyzing a number of small and a little larger examples.




Figure 4.14. An example of a MOP

Project 6

Assume a graph consists of a cycle, and three more vertices, each adjacent to just one other vertex on the cycle. What can be said about the Nash equilibria in doctor location games on such graphs?

Project 7

Is it possible to determine which symmetric simultaneous constant-sum games with constant sum of payoffs n equal to the number of moves of each player are doctor location games?

CHAPTER 5

Example: Restaurant Location Games

Prerequisites: Chapters 1 and 2.

As in the previous chapter we discuss games on graphs. Their most interesting feature is that they always have pure Nash equilibria.

Let us repeat the definition of an undirected graph, given in Chapter 4. Undirected **graphs** have **vertices**, displayed by small circles. Some pairs of vertices are connected by undirected **edges**. They are drawn as curves connecting the two vertices. Vertices connected by an edge are called **adjacent**. Adjacent vertices are also called **neighbors**.

Again we use graphs to model islands. The vertices represent the towns, and the edges the roads between the towns.

Restaurant Location Game On a small island are towns, all having the same number of residents. The towns are connected by a system of streets. Two sushi restaurants, Ann's and Beth's, are about to open. Both have the same quality and price level and are almost indistinguishable. Market research shows:

- People with a sushi restaurant in their home town would on average eat sushi once a year. They would not visit sushi restaurants outside their home town. If both sushi restaurants were in the home town, residents would randomly choose either, or alternate their choices.
- A person with no sushi restaurant in the home town, but with a restaurant in an adjacent town, would on average eat sushi once in two years. If more than one sushi restaurants were in a town adjacent to the person's home town, the person would randomly choose either, or alternate choices.
- Nobody would go to a sushi restaurant that is not in the home town or an adjacent town.

In the simultaneous version of the game, both Ann and Beth simultaneously submit a proposal for the location of their restaurants. The payoffs for the players are the expected numbers of customers per year. What location should the players choose?

Student Activity Play the game in one of the applets [Location1](#), [Location2](#), [Location6](#), [Location10](#), or [Location11](#). In the applets [LocationGrid8](#) and [LocationGrid16](#) you can create your own graph.

Unlike the doctor version, the restaurant version is not a constant-sum game, since the total number of customers varies. If the restaurants are placed far apart, the total number of customers would normally be

higher than if they are located in adjacent towns or in the same town. If the graph is large, then the restaurants might find their own niches far apart from one another, and there would not be much competition between them. This is different from the doctor-location games. The game is more interesting if the graph is smaller, when some competition between the restaurants is unavoidable.

In his book *The Wealth of Nations* Adam Smith coined the phrase “the invisible hand”. This phrase stands for the claim that, if government allows agents to act freely, a solution best for the whole society will occur. Some mechanisms work like this and others don’t. We have seen in Chapter 4 that letting doctors choose their location does not necessarily result in a solution optimal for the people. In the restaurant location model, at least in the examples we consider here, we will see that the invisible hand seems to work, if not perfectly.

How are the payoffs computed in the three cases where the restaurants are placed in the same town, in adjacent towns, or in nonadjacent towns? It turns out that, under the assumption that all towns have about the same size, they depend only on three quantities: the number $d(x)$ of neighbors of vertex x , the number $d(y)$ of neighbors of y , and the number $n(x, y)$ of common neighbors of the vertices x and y , where x and y are the locations of Ann’s and Beth’s restaurants.

- If both Ann’s and Beth’s restaurants are placed in the same town x , then they share the residents of x as customers. In addition, they share the residents of the adjacent towns every two years, meaning their customer share each year is half the sum of the sizes of adjacent towns. Therefore each restaurant can expect $\frac{1}{2} + \frac{1}{4} \cdot d(x)$ towns as customers per year.
- If the restaurants are placed in different nonadjacent towns x and y , then x contributes 1 to Ann’s payoff, each of the $n(x, y)$ common neighbors of x and y adds $\frac{1}{4}$, and each of the $d(x) - n(x, y)$ other neighbors of x contributes $\frac{1}{2}$ to Ann’s payoff. Thus Ann’s payoff is

$$1 + \frac{1}{2} \cdot (d(x) - n(x, y)) + \frac{1}{4} \cdot n(x, y) = 1 + \frac{1}{2} \cdot d(x) - \frac{1}{4} \cdot n(x, y)$$

and Beth’s payoff is

$$1 + \frac{1}{2} \cdot (d(y) - n(x, y)) + \frac{1}{4} \cdot n(x, y) = 1 + \frac{1}{2} \cdot d(y) - \frac{1}{4} \cdot n(x, y).$$

- If the restaurants are placed in adjacent towns x and y , then Ann gets all inhabitants from x , $\frac{1}{4}$ from each one of the $n(x, y)$ towns adjacent to both x and y , and $\frac{1}{2}$ from each other neighbor of x except y . Therefore Ann’s payoff is

$$1 + \frac{1}{2} \cdot (d(x) - 1 - n(x, y)) + \frac{1}{4} \cdot n(x, y) = \frac{1}{2} + \frac{1}{2} \cdot d(x) - \frac{1}{4} \cdot n(x, y)$$

and Beth’s payoff is

$$1 + \frac{1}{2} \cdot (d(y) - 1 - n(x, y)) + \frac{1}{4} \cdot n(x, y) = \frac{1}{2} + \frac{1}{2} \cdot d(y) - \frac{1}{4} \cdot n(x, y).$$

5.1 A First Graph

Look at the graph in Figure 5.1. You can play restaurant placement in the applet [Location6](#). Let me explain in three examples how the payoff matrix of the game is derived.

- If Ann chooses town 1 and Beth town 2, Ann’s restaurant attracts all people from town 1, one fourth of town 3, and one half of town 5 per year, making an expected payoff of $\frac{7}{4}$ towns per year. Beth’s restaurant attracts all people from town 2, one fourth of town 3, and one half of town 7, adding to the same expected payoff. We could also use the formulas above. x and y are adjacent, and $d(x) = d(y) = 3$ and $n(x, y) = 1$. Then the payoff for each is $\frac{1}{2} + \frac{1}{2} \cdot 3 - \frac{1}{4} \cdot 1 = \frac{7}{4}$.




Figure 5.1. A graph

- If Ann chooses town 2 and Beth town 5, then we can use the formula for nonadjacent locations. Again $d(x) = d(y) = 3$ and $n(x, y) = 1$, so Ann’s and Beth’s payoffs are $1 + \frac{1}{2} \cdot 3 - \frac{1}{4} \cdot 1 = \frac{9}{4}$.
- If both choose town 1, then we can use the formula with $d(x) = 3$ to get payoffs of $\frac{1}{2} + \frac{1}{4} \cdot 3 = \frac{5}{4}$ for both Ann and Beth.

Doing the computations, we arrive at the payoff bimatrix

	1	2	3	4	5	6	7
1	1.25, 1.25	1.75, 1.75	1.75, 1.75	2, 1.5	2, 2	2.25, 1.75	2.25, 1.75
2	1.75, 1.75	1.25, 1.25	1.75, 1.75	2.25, 1.75	<u>2.25, 2.25</u>	2.25, 1.75	2, 1.5
3	1.75, 1.75	1.75, 1.75	1.25, 1.25	2, 1.5	2, 2	<u>2.5, 2</u>	2.25, 1.75
4	1.5, 2	1.75, 2.25	1.5, 2	1, 1	1.5, 2	1.75, 1.75	2, 2
5	2, 2	<u>2.25, 2.25</u>	2, 2	2, 1.5	1.25, 1.25	2, 1.5	2.25, 1.75
6	1.75, 2.25	1.75, 2.25	<u>2, 2.5</u>	1.75, 1.75	1.5, 2	1, 1	1.5, 1.5
7	1.75, 2.25	1.5, 2	1.75, 2.25	2, 2	1.75, 2.25	1.5, 1.5	1, 1

The maximin moves for both players are 1, 2, 3, and 5, exactly the vertices with three neighbors. The security level achieved by these moves is $\frac{5}{4}$.

There are four Nash equilibria, underlined in the table. Two of them are shown in Figure 5.2; the other two are obtained by interchanging Ann’s and Beth’s moves.




Figure 5.2. Two Nash equilibria

For this graph they are the only outcomes where every town is adjacent to a restaurant, but of course this is not true in general. Still, the four Nash equilibria are the only outcomes where the sum of the payoffs reaches the maximum value of 4.5.

5.2 A Second Graph

Play the game in the applet [Location11](#).




Figure 5.3. The second graph

	1	2	3	4	5	6	7	8
1	1.5, 1.5	2.25, 1.75	2.5, 2	2.25, 1.25	<u>2, 2.5</u>	2, 2	2.5, 2	3, 2
2	1.75, 2.25	1.25, 1.25	1.75, 1.75	2.25, 1.75	1.5, 2.5	2, 2.5	2.25, 2.25	2.25, 1.75
3	2, 2.5	1.75, 1.75	1.25, 1.25	2.5, 2	1.75, 2.75	2.25, 2.75	2, 2	2, 1.5
4	1.25, 2.25	1.75, 2.25	2, 2.5	1, 1	1.5, 3	1.25, 2.25	1.75, 2.25	2, 2
5	<u>2.5, 2</u>	2.5, 1.5	2.75, 1.75	3, 1.5	1.75, 1.75	<u>2.5, 2</u>	2.75, 1.75	3, 1.5
6	2, 2	2.5, 2	2.75, 2.25	2.25, 1.25	<u>2, 2.5</u>	1.5, 1.5	2.25, 1.75	2.75, 1.75
7	2, 2.5	2.25, 2.25	2, 2	2.25, 1.75	1.75, 2.75	1.75, 2.25	1.25, 1.25	2, 1.5
8	2, 3	1.75, 2.25	1.5, 2	2, 2	1.5, 3	1.75, 2.75	1.5, 2	1, 1

This game has four pure Nash equilibria, underlined in the bimatrix. In both the sum of payoffs is 4.5. However, the combination of one restaurant in town 1 and the other one in town 8 serves the community better, since the sum of payoffs is 5, maximizing the number of inhabitants reached.

Move number 5 is the only maximin move, with a security level of 1.75. The maximum degree of the graph is $\Delta = 5$. Vertex 5 is the vertex whose degree equals the maximum degree Δ . This is not a coincidence:

Fact For every graph, the maximin moves are locations on vertices with maximum degree $d(x) = \Delta$. The security level is $\frac{1}{2} + \frac{\Delta}{4}$.

Why? If Ann plays vertex x , the worst that could happen to Ann is when Beth also places her restaurant on vertex x , since both payoffs then are $\frac{1}{2} + \frac{d(x)}{4}$, which is always smaller or equal to the Ann's payoffs $1 + \frac{d(x)}{2} - \frac{n(x,y)}{4}$ respectively $\frac{1}{2} + \frac{d(x)}{2} - \frac{n(x,y)}{4}$ when Beth places her restaurant somewhere else (since $n(x, y) \leq d(x)$). Thus in row x , the minimum value for Ann is $\frac{1}{2} + \frac{d(x)}{4}$. This value is maximized for x with maximal $d(x)$.

5.3 Existence of Pure Nash Equilibria

For every vertex x , the payoff Ann would get if she locates the restaurant in x and Beth places her restaurant far away is $1 + \frac{d(x)}{2}$. The same is true for Beth, if she places her restaurant at y and Ann places far away: Then Beth's payoff is $1 + \frac{d(y)}{2}$. But if the locations x and y are closer together, both Ann's and Beth's payoffs are reduced. Moreover, and this is the important part, both Ann's and Beth's payoffs are reduced by the same amount,

$$g(x, y) = \begin{cases} \frac{1}{2} + \frac{d(x)}{4} & \text{if } x \text{ and } y \text{ are identical} \\ \frac{1}{2} + \frac{n(x,y)}{4} & \text{if } x \text{ and } y \text{ are distinct and adjacent} \\ \frac{n(x,y)}{4} & \text{if } x \text{ and } y \text{ are distinct and nonadjacent.} \end{cases}$$

Theorem [P2011] *All simultaneous restaurant location games with one restaurant for each player have at least one Nash equilibrium in pure strategies.*

We can describe how to find a Nash equilibrium. Starting with a vertex, we proceed to one of its best responses, and from there to one of its best responses, and so on: we start with a vertex and move along arcs of the best response digraph. Eventually we must visit a vertex that we have seen already in the sequence, so we will find a directed cycle in the best response digraph. We will discuss only one special case, but others can be treated similarly.

The special case is where the cycle consists of four vertices x, y, z , and w , i.e., where y is a best response to x , z a best response to y , w is a best response to z , and x is a best response to w . Since y is the best response to x , and therefore at least as good a response as w , we get

$$1 + \frac{d(y)}{2} - g(y, x) \geq 1 + \frac{d(w)}{2} - g(w, x).$$

In the same way, we have

$$\begin{aligned} 1 + \frac{d(z)}{2} - g(z, y) &\geq 1 + \frac{d(x)}{2} - g(x, y) \\ 1 + \frac{d(w)}{2} - g(w, z) &\geq 1 + \frac{d(y)}{2} - g(y, z) \\ 1 + \frac{d(x)}{2} - g(x, w) &\geq 1 + \frac{d(z)}{2} - g(z, w). \end{aligned}$$

The sum of the four left sides of the inequalities is greater or equal to the sum of the four right sides:

$$\begin{aligned} 1 + \frac{d(y)}{2} - g(y, x) + 1 + \frac{d(z)}{2} - g(z, y) + 1 + \frac{d(w)}{2} - g(w, z) + 1 + \frac{d(x)}{2} - g(x, w) \\ \geq 1 + \frac{d(w)}{2} - g(w, x) + 1 + \frac{d(x)}{2} - g(x, y) + 1 + \frac{d(y)}{2} - g(y, z) + 1 + \frac{d(z)}{2} - g(z, w). \end{aligned}$$

But since $g(x, y) = g(y, x)$, and so on, the terms on the left are exactly the ones on the right, so we have equality. That means that there was equality in all inequalities. Equality in the first implies that w and y are best responses to x , and since x is a best response to w we have a Nash equilibrium x versus w . The other three equations give three more pairs of Nash equilibria— x versus y , y versus z , and z versus w . So we have found Nash equilibria.

5.4 More than one Restaurant (optional)

We could play variants where two restaurant chains place several restaurants each. Because we are playing a simultaneous game, all location decisions are made simultaneously by the two players. If a town hosts no restaurant, but if two restaurants from Ann's chain and one from Beth's chain are adjacent to the town,

then half of the population would go to a restaurant, one fourth to one from Ann's chain and one fourth to Beth's restaurant. Thus they would not visit every adjacent restaurant with the same frequency. This could be explained by 50% of the population of every town having a slight preference towards Ann's chain, and 50% towards Beth's chain, so slight that they would always go to the closer one no matter what brand, but in case of a tie of distance, they would go to their preferred brand.

Let us discuss what can happen at the example of Graph 10 in Figure 5.4. Play the game in the applet [Location10](#). Since you cannot click simultaneously with one mouse, in the applet you have to select Ann's first restaurant, Beth's first restaurant, Ann's second restaurant, and Beth's second restaurant, in that order.




Figure 5.4. Graph 10

Let us abbreviate a move where a player places her two restaurants at vertices x and y as $\{x, y\}$. Since every player has the 9 moves $\{1, 1\}, \dots, \{9, 9\}$ and the 36 moves $\{x, y\}$ with x and y different ($\{x, y\}$ and $\{y, x\}$ are the same), each player has 45 possible moves, so the payoff bimatrix would have dimensions 45×45 —too large for us to display and discuss.

Still we can check that the method of finding Nash equilibria inside directed cycles of the best response digraph for restaurant location games with one restaurant each fails. Assume Ann places her restaurant on 4 and 6, i.e., chooses move $\{4, 6\}$. Confirm that one of Beth's best responses to that would be $\{2, 5\}$ (the other best response would be $\{1, 2\}$). Ann's best response to $\{2, 5\}$ is $\{1, 9\}$ (or $\{4, 9\}$). Beth's best response to $\{1, 9\}$ is $\{4, 6\}$. So we have a directed cycle of length three in the best response digraph, but unlike the one restaurant case, the reverse arcs are absent: $\{1, 9\}$ is not a best response to $\{4, 6\}$, $\{2, 5\}$ is not a best response to $\{1, 9\}$, and $\{4, 6\}$ is not a best response to $\{2, 5\}$.

Looking at the whole best response digraph, we can confirm that the game has no pure Nash equilibria.

Exercises

1. Look at a simultaneous restaurant location game with two restaurants for each player. Find moves $\{x, z\}$ for Ann and $\{y, w\}$ for Beth where the difference between Ann's actual payoff and the payoff she would get if both of Beth's restaurants were far away is different from the difference between Beth's actual payoff and the payoff she would get if both of Ann's restaurants were far away. In the 1-restaurant case, both these numbers must be equal (denoted $g(x, y)$).
2. Look at a variant of the 1-restaurant location game where the average person would visit the closest restaurant once a year if it is in the home town, every other year if it is in a neighbor town, and every four years if it is in a neighbor town of a neighbor town. Look at moves 7 and 8 in Graph 10 in Figure 5.4. Calculate the payoffs for the players for that case. What would the payoffs be if the placement of the other player were far away, assuming that the play graph were much larger than Graph 10? Which one of the players is hurt more by the placement of the other player?




Figure 5.5. Four graphs

3. The bimatrix is the bimatrix of graph A in Figure 5.5:

	1	2	3	4	5	6	7	8
1	1.5, 1.5	—, —	—, —	2.25, 1.25	2, 2.5	2, 2	2.5, 2	3, 2
2	1.75, 2.25	1.25, 1.25	1.75, 1.75	—, —	1.5, 2.5	2, 2.5	2.25, 2.25	2.25, 1.75
3	2, 2.5	1.75, 1.75	1.25, 1.25	—, —	1.75, 2.75	2.25, 2.75	2, 2	2, 1.5
4	1.25, 2.25	—, —	—, —	1, 1	1.5, 3	1.25, 2.25	1.75, 2.25	2, 2
5	2.5, 2	2.5, 1.5	2.75, 1.75	3, 1.5	—, —	2.5, 2	2.75, 1.75	3, 1.5
6	2, 2	2.5, 2	2.75, 2.25	2.25, 1.25	2, 2.5	1.5, 1.5	2.25, 1.75	2.75, 1.75
7	2, 2.5	2.25, 2.25	2, 2	2.25, 1.75	1.75, 2.75	1.75, 2.25	1.25, 1.25	2, 1.5
8	2, 3	1.75, 2.25	1.5, 2	2, 2	1.5, 3	1.75, 2.75	1.5, 2	1, 1

Find the missing values, and find all pure Nash equilibria.

4. The bimatrix is the bimatrix of graph B in Figure 5.5.

	1	2	3	4	5	6	7	8
1	1.5, 1.5	2.25, 1.75	2.5, 2	2.25, 1.25	2, 3	2, 2	2.5, 2	2.75, 2.25
2	—, —	1.25, 1.25	—, —	—, —	1.5, 3	2, 2.5	2.25, 2.25	2, 2
3	—, —	—, —	1.25, 1.25	2.5, 2	1.5, 3	2.25, 2.75	2, 2	1.75, 1.75
4	—, —	—, —	2, 2.5	1, 1	1.5, 3.5	1.25, 2.25	1.75, 2.25	2, 2.5
5	3, 2	3, 1.5	3, 1.5	3.5, 1.5	2, 2	3, 2	3, 1.5	3, 1.5
6	2, 2	2.5, 2	2.75, 2.25	2.25, 1.25	2, 3	—, —	2.25, 1.75	2.5, 2
7	2, 2.5	2.25, 2.25	2, 2	2.25, 1.75	1.5, 3	1.75, 2.25	1.25, 1.25	1.75, 1.75
8	2.25, 2.75	2, 2	1.75, 1.75	2.5, 2	1.5, 3	2, 2.5	1.75, 1.75	1.25, 1.25

Find the missing values, and find all pure Nash equilibria.

5. The bimatrix is the bimatrix of graph C in Figure 5.5.

	1	2	3	4	5	6	7	8
1	1.5, 1.5	2.25, 1.75	2.5, 2	2, 1.5	1.75, 2.75	—, —	—, —	—, —
2	1.75, 2.25	1.25, 1.25	1.75, 1.75	2, 2	1.5, 3	2, 2.5	2.5, 2	2, 2
3	2, 2.5	1.75, 1.75	1.25, 1.25	2.25, 2.25	—, —	2.25, 2.75	2.25, 1.75	1.75, 1.75
4	1.5, 2	2, 2	2.25, 2.25	1.25, 1.25	—, —	1.5, 2	2.25, 1.75	2.25, 2.25
5	2.75, 1.75	3, 1.5	—, —	—, —	2, 2	3, 2	3.5, 1.5	3.25, 1.75
6	2, 2	2.5, 2	2.75, 2.25	2, 1.5	2, 3	—, —	2.5, 1.5	2.5, 2
7	1.75, 2.75	2, 2.5	1.75, 2.25	1.75, 2.25	1.5, 3.5	1.5, 2.5	1, 1	1.5, 2
8	2.25, 2.75	2, 2	1.75, 1.75	2.25, 2.25	1.75, 3.25	2, 2.5	2, 1.5	1.25, 1.25

Find the missing values, and find all pure Nash equilibria.

6. The bimatrix is the bimatrix of graph D in Figure 5.5.

	1	2	3	4	5	6	7	8
1	1.25, 1.25	1.75, 1.75	—, —	1.5, 2	—, —	1.75, 2.25	2.25, 2.25	—, —
2	1.75, 1.75	1.25, 1.25	1.75, 2.25	1.5, 2	2.25, 2.75	2, 2.5	2.5, 2.5	2.25, 2.25
3	—, —	2.25, 1.75	1.5, 1.5	2.25, 2.25	2.25, 2.25	2.5, 2.5	2.5, 2	—, —
4	2, 1.5	2, 1.5	2.25, 2.25	—, —	2.5, 2.5	2.25, 2.25	2.75, 2.25	2.75, 2.25
5	—, —	2.75, 2.25	2.25, 2.25	2.5, 2.5	1.5, 1.5	2.25, 2.25	2, 1.5	2, 1.5
6	2.25, 1.75	2.5, 2	2.5, 2.5	2.25, 2.25	2.25, 2.25	1.5, 1.5	2.25, 1.75	2.5, 2
7	2.25, 2.25	2.5, 2.5	2, 2.5	2.25, 2.75	1.5, 2	1.75, 2.25	1.25, 1.25	1.75, 1.75
8	2.5, 2.5	2.25, 2.25	1.75, 2.25	2.25, 2.75	1.5, 2	2, 2.5	1.75, 1.75	1.25, 1.25

Find the missing values, and find all pure Nash equilibria.

7. With [LocationGrid8](#) create your own random graph by clicking the "start new graph" button. Find a Nash equilibrium without creating the bimatrix—use the method of starting somewhere, finding the best response to this move, and finding the best response to the best response, and so on, until you have found a Nash equilibrium. Submit the graph and the Nash equilibrium pair.
8. Do the same with applet [LocationGrid16](#).

CHAPTER 6

Using Excel

During my senior year in high school I became interested in game theory. At that time I also learned my first programming language, FORTRAN, in a summer course at the University of Stuttgart. The first game I analyzed with computer help was simultaneous QUATRO-UNO described in Project 8. Since then I have believed that analyzing games formally may give new insights only if the game is sufficiently complex, and that for analyzing complex games technology is required. As evidence, this book is exhibit A.

Let's move from my private belief to the history of science and technology. Von Neumann and Morgenstern's book on game theory was published in 1944. A few years later, game theory attracted a lot of attention at Princeton University and at the RAND corporation. On the other hand, the first working computers appeared a few years earlier. A coincidence? Maybe, but if computers were not invented, von Neumann's ideas might have gone unnoticed for years or decades.

Game theory, as we have seen for two-player simultaneous games and others, works with large bimatrices (tables). So far the operations required to analyze a game—maximin moves, domination, best responses and pure Nash equilibria—are simple, but tedious, so computer help could be useful. To analyze these games, why not use spreadsheet programs, designed to work with tables?

This chapter gives a brief introduction to Excel, a spreadsheet program, and provides several Excel files that can be used to analyze simultaneous games if the (bi-) matrices are given. You are supposed to use the files for larger examples from now on. Later, we will provide other Excel sheets that help to create (bi-) matrices.

6.1 Spreadsheet Programs like Excel

Excel, part of Microsoft Office, is probably the most popular spreadsheet program, but others like Open Office could also be used, except for a few that contain macros. **Spreadsheet programs** perform calculations on data that is grouped in tables, with rows and columns. Rows are named by letters, and columns by numbers. **Cells** are labeled by column letter and row number, like "C5", and contain either data—numbers or text—or formulas. A formula refers to other cells and calculates a value based on the values in them. For instance, if we put a value of 8 into cell A1 and a value of 9 into cell A2, and the formula " $=A1+A2$ " into cell A3, then cell A3 will show the sum of the values in A1 and A2, in our case 17. If you change the 8 in cell A1 to 7, the 17 in cell A3 will automatically update to 16. A formula can even refer to a cell that also contains a formula, in which case instead of the data in the cell the displayed value of the cell is taken into the calculation.

If you don't want to create Excel sheets but only use them, this is about all you need to know. You are supposed to input some data into some cells, the input cells. Based on the data, you will get output as values in output cells where functions have been written in. You have to be a little cautious with the output values shown, since they may be rounded. But even if the display of a value is rounded, formulas based on it will always use the precise number.

In addition to output in output cells, sometimes you will see some formatting based on the input. Output cell values may be displayed in charts, to allow you to grasp data faster.

Before we move into examples, a final hint: an Excel file often contains different pages, that can be accessed on tabs at the bottom of the tables. There may be connections between the pages, formulas on one page referring to cells on other pages.

6.2 Two-Person Simultaneous Games

The Excel file [SIMnonzero.xlsx](#) should be used for general simultaneous 2-player games, where each player has up to 21 options. Open the file and start at the page called “data” by clicking the corresponding tab below to the left. In the blue table, write the payoffs for both players in the corresponding pairs of cells. If Ann and Beth have fewer than 21 possible moves, leave the corresponding rows or columns empty.

On the four pages called “maximin”, “weak dom”, “best response”, and “pure Nash”, the analysis is done.

- On “maximin”, the maximin moves for Ann and Beth are highlighted in red. The security level is displayed in orange either on the bottom (for Beth) or to the right of the table (for Ann).
- On the page “weak dom”, weakly dominated moves are shown in gray.
- On the page “best response”, the best responses are highlighted in green.
- On the page “pure Nash”, the pairs of payoffs attached to pure Nash equilibria are highlighted in green.

For zero-sum games use the file [SIMzero.xlsx](#), where you don’t have to type in Beth’s payoffs. Since in each outcome Beth’s payoff is the negative of Ann’s payoff, Beth’s payoffs are generated automatically on the page “non-zero-sum” when you type Ann’s payoffs in the table on the page called “data”.

6.3 Three-Person Simultaneous Games

For three players we use the Excel sheets [SIM3.xlsx](#) or [SIM33.xlsx](#), depending on whether each player has two or three options. Each file contains three pages. Start on the page called “data”. You are supposed to fill in numbers, payoffs for Ann, Beth, and Cindy in the corresponding outcome, into the blue cells.

On the tab labeled “maximin”, the maximin moves are highlighted in orange.

The tab called “pure Nash” displays the pure Nash equilibria outcomes in two shades of green. The computation can be seen on the lower part of the sheet. For each outcome, we compute Ann’s, Beth’s, and Cindy’s best response. The words “no” and “yes” tell whether the best response of the player is the same as the option chosen by the player. Only if this is true, i.e., if the word “yes” appears three times in a row, we have a Nash equilibrium.

Exercises

1. Analyze the 2-person game with [SIMnonzero.xlsx](#):

	B_1	B_2	B_3	B_4
A_1	1,3	4,2	1,1	3,5
A_2	2,4	2,5	3,5	4,1
A_3	3,1	2,2	1,2	2,3
A_4	3,2	2,1	3,1	2,4
A_5	4,2	3,3	1,5	3,2

2. Analyze the 7×7 simultaneous 2-person game whose payoff bimatrix can be found on the “Exercise 2” tab in the Excel file [Exercises.xlsx](#).

3. Analyze the 9×9 simultaneous 2-person game whose payoff bimatrix can be found on the “Exercise 3” tab in [Exercises.xlsx](#).
4. Analyze the 10×10 simultaneous 2-person game whose payoff bimatrix can be found on the “Exercise 4” tab in [Exercises.xlsx](#).
5. Analyze the 6×6 simultaneous 2-person game whose payoff bimatrix can be found on the “Exercise 5” tab in [Exercises.xlsx](#).
6. Analyze the 10×10 simultaneous zero-sum 2-person game whose payoff matrix can be found on the “Exercise 6” tab in [Exercises.xlsx](#).
7. Analyze the 10×10 simultaneous zero-sum 2-person game whose payoff matrix can be found on the “Exercise 7” tab in [Exercises.xlsx](#).
8. Analyze the 8×8 simultaneous zero-sum 2-person game whose payoff matrix can be found on the “Exercise 8” tab in [Exercises.xlsx](#).

Project 8

Simultaneous QUATRO-UNO There are two players, and each gets four cards of values 1, 2, 3, and 4. Each player orders her four cards into a pile, which is then put face-down on the desk. Then repeatedly both players reveal their topmost card. The lower-valued card is taken away from the pile, and the higher one, the winning one, remains on the pile. In case of identical cards, both cards are removed. There is also one exception: if one card is a “1” and the other one a “4”, then too both cards are removed. After each round, the player whose topmost card was removed reveals the next card. The first player with an empty pile loses.

Play ten games against the computer in the applet [QuatroUnosimCom](#) or ten games against a friend in applet [QuatroUnosim](#). The payoff matrix is given on the “Exercise 10” tab in the file [Exercises.xlsx](#). Use the second applet to replace the five question marks by values.

Try to find the IEWD matrix—what remains after repeatedly eliminating weakly dominated moves. Since the matrix is 24×24 , too large for our Excel file [SIMzero.xlsx](#), you have to find three weakly dominated moves by hand. Compare 4123 and 4132, for instance. Then compare 1234 and 1324. Maybe you see a pattern and can find a third pair of moves where one weakly dominates the other.

Are there any pure Nash equilibria? This project will be continued as Project 53 in Chapter 27.

Project 9

Use [LocationGame.xlsm](#) to compute all Nash equilibria of the Restaurant Location Game for several graphs, and test how many of them maximize the sum of the payoffs of the two players. If they don’t, what is the worst ratio of sum of player’s payoffs in a Nash equilibrium and the maximum possible sum of player’s payoffs?

The next two projects refer to the following related games. The 5 KNIGHTS game has been discussed in Chapter 2, see the description there. You can experiment and search for Nash equilibria in the applets [5Knights](#) and [5ExhaustedKnights](#), and in [5ExhaustedKnightsRand](#) for randomly generated preferences.

5 CARDINALS Five cardinals, A, B, C, D, E, are electing their leader. Each has a list of preferences, for example (preferences given from highest to lowest):

A: A, B, C, D, E

B: B, C, A, E, D

C: E, C, A, D, B

D: C, E, B, D, A

E: E, C, B, A, D

They elect in rounds. In each round, cardinals simultaneously submit one name each. Since cardinals are more concerned about majorities than knights, a cardinal is elected only if he gets three or more votes. If none gets more than two votes, we proceed to the next round. There are two versions:

Early Case: If the cardinal's first choice is elected, his payoff is 2. If the cardinal's second choice is elected, his payoff is 1. If nobody is elected, the payoff is 0. If the cardinal's third, fourth, or fifth choice is elected, the cardinal's payoff is -1 , -2 , or -3 .

Exhausted Case: The cardinal's first, second, or third choice gives a payoff of 2, 1, 0. If no one is elected, the payoff is -1 . If the cardinal's fourth or fifth choice is elected, the payoff is -2 or -3 .

You can experiment and search for Nash equilibria in the applets [5Cardinals](#) and [5ExhaustedCardinals](#).

Project 10

For this project, you should use the Excel sheet [5Knights.xlsx](#). It calculates all Nash equilibria for both versions for any preferences. Discuss the following questions:

- Which version has in general more Nash equilibria—the early case version or the exhausted case version? Which version has more Nash equilibria that are not ties?
- Can you predict, by looking at the preferences of the five players, which player is elected most times in Nash equilibria?
- Are there always Nash equilibria in the early case version that are not Nash equilibria in the exhausted case version? Can you describe them? Are there always Nash equilibria in the exhausted case version that are not Nash equilibria in the early case version? Can you describe them?

You can experiment in the applets [5Knights](#), [5ExhaustedKnights](#), and also [5ExhaustedKnightsRand](#) for random preferences.

Project 11

For this project, you should use the Excel sheet [5Cardinals.xlsx](#), and in (a) also [5Knights.xlsx](#). It calculates all Nash equilibria for both versions for any preferences. Discuss the following questions:

- Which version has in general more Nash equilibria—the early case version or the exhausted case version? Which version has more Nash equilibria that are not ties? What about the numbers compared to the two versions of the 5 KNIGHT game with same preferences?
- Can you predict, by looking at the preferences of the five players, which player is elected most times in Nash equilibria in the 5 CARDINALS game?

- c) For the 5 CARDINALS games, are there always Nash equilibria in the early case version that are not Nash equilibria in the exhausted case version? Can you describe them? Are there always Nash equilibria in the exhausted case version that are not Nash equilibria in the early case version? Can you describe them?

CHAPTER 7

Example: Election I

Prerequisites: Chapters 1, 2, and 6.

The president of the USA is elected by electors from all 50 states. All the electoral votes from a state go to the most popular candidate in that state. If one week before the election a candidate knows that she is behind in a few states, and leading in others, what would be a good strategy for the remaining time? Concentrating on those states where she is behind, or accepting that they are lost and concentrating on others? The decision will depend on several factors, including whether the state can still be turned, and on the size of the state. California is more important than Montana in the presidential election. In this chapter we look at a simplified model with only three districts and analyze a few cases formally. Looking at simpler cases may allow us to extract rules for larger situations.

7.1 First Example

We start with a special game, part of a family of games that will be introduced later.

ELECTION 1 or ELECTION(7, 8, 13 | -1, -1, 1 | 3, 3) In Kalumba there are three electoral districts, C, D, and E. As in the election of the President of the USA, the President of Kalumba is elected by electoral votes. There are 7 electoral votes from district C, 8 from district D, and 13 from district E. Districts do not split electoral votes. There are two presidential candidates, Ann and Beth, and in the last phases of their campaigns they simultaneously decide how to allocate the three remaining resources each has. Each must be allocated entirely to one district. A district may receive more than one resource. Each district will vote for the candidate who put more resources into the district (not just during the last phase), and will abstain in case of a tie. In districts C and D, Ann is 1 resource unit behind, and in district E, Ann has an advantage of 1 resource unit. How should Ann and Beth distribute their resources?

There are three outcomes, win, loss, draw, so we will model this as a zero-sum game with payoffs 1 (win), 0 (draw), and -1 (lose). We have a 2-player simultaneous game with many moves. Let us first list all moves:

A move means allocating x resources to district C, y resources to district D, and z resources to district E, obeying the resource restriction $x + y + z = 3$. This is abbreviated as (x, y, z) . Furthermore x , y , and z must be nonnegative integers. Thus Ann and Beth have the ten moves (0, 0, 3), (1, 0, 2), (0, 1, 2), (2, 0, 1), (1, 1, 1), (0, 2, 1), (3, 0, 0), (2, 1, 0), (1, 2, 0), and (0, 3, 0).

Calculating the outcome and therefore the payoffs for a pair of moves is easy. Let's demonstrate this at the example where Ann plays (0, 1, 2) and Beth plays (1, 0, 2). Beth puts one more effort into district C, and already had the lead there, so Beth wins district C. Ann puts one more resource into district D, in the last

phase, but since Beth was leading there by 1, we get a tie in district D. For district E, both put the same effort in the last phase, but since Ann had an advantage of 1 before the last phase, Ann wins district E. Therefore Ann gets 13 votes (from district E), and Beth gets 7 votes (from district C). District D is undecided and doesn't vote for either candidate. Thus Ann is elected president for that pair of moves, and gets a payoff of 1, while Beth's payoff is -1 .

Student Activity Before proceeding with the mathematical analysis, think how Beth should play.

Ann's payoffs for the $10 \cdot 10$ pairs of moves are given in the matrix below:

	0,0,3	1,0,2	0,1,2	2,0,1	1,1,1	0,2,1	3,0,0	2,1,0	1,2,0	0,3,0
0,0,3	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
1,0,2	-1	-1	1	-1	-1	1	-1	-1	-1	1
0,1,2	-1	1	-1	1	-1	-1	1	-1	-1	-1
2,0,1	-1	-1	-1	-1	1	1	-1	-1	1	1
1,1,1	-1	-1	-1	1	-1	1	1	-1	-1	1
0,2,1	-1	1	-1	1	1	-1	1	1	-1	-1
3,0,0	-1	-1	-1	-1	-1	-1	-1	1	1	1
2,1,0	-1	-1	-1	-1	-1	-1	1	-1	1	1
1,2,0	-1	-1	-1	1	-1	-1	1	1	-1	1
0,3,0	-1	-1	-1	1	1	-1	1	1	1	-1

If Beth plays $(0, 0, 3)$ she wins, no matter what Ann does. The explanation is simple. If Ann also invests all three resources into district E, then Ann keeps E, but Beth keeps C and D where she was leading before, and since $7 + 8 > 13$, Beth wins. If Ann puts two resources into district E, then there is a draw in E and in the other district where Ann invests. But Beth keeps the third district and wins. Finally, if Ann invests less than two resources into district E, then Beth wins district E, but Ann cannot win (turn around) both districts C and D with the three resources she has, so Beth wins then as well.

7.2 Second Example

ELECTION 2 or ELECTION(7, 8, 13 | -1, -1, 1 | 4, 4) We modify ELECTION 1 by allowing four resources for each player

We have even more moves here. Ann and Beth have the $1 + 2 + 3 + 4 + 5 = 15$ moves $(0, 0, 4)$, $(1, 0, 3)$, $(0, 1, 3)$, $(2, 0, 2)$, $(1, 1, 2)$, $(0, 2, 2)$, $(3, 0, 1)$, $(2, 1, 1)$, $(1, 2, 1)$, $(0, 3, 1)$, $(4, 0, 0)$, $(3, 1, 0)$, $(2, 2, 0)$, $(1, 3, 0)$, and $(0, 4, 0)$. I listed first the move with $z = 4$, then the two moves with $z = 3$, and so on. Beth can no longer always win by putting all her resources into district E, since Ann could respond by putting two resources into district C and D each, thereby winning both.

The 15×15 matrix is best generated by a computer. We can do it in Excel sheet [Election.xlsx](#). The payoff matrix is

	0,0,4	1,0,3	0,1,3	2,0,2	1,1,2	0,2,2	3,0,1	2,1,1	1,2,1	0,3,1	4,0,0	3,1,0	2,2,0	1,3,0	0,4,0
0,0,4	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
1,0,3	-1	-1	1	-1	-1	1	-1	-1	-1	1	-1	-1	-1	-1	1
0,1,3	-1	1	-1	1	-1	-1	1	-1	-1	-1	1	-1	-1	-1	-1
2,0,2	-1	-1	-1	-1	1	1	-1	-1	1	1	-1	-1	-1	1	1
1,1,2	-1	-1	-1	1	-1	1	1	-1	-1	1	1	-1	-1	-1	1
0,2,2	-1	1	-1	1	1	-1	1	1	-1	-1	1	1	-1	-1	-1
3,0,1	-1	-1	-1	-1	-1	-1	-1	1	1	1	-1	-1	1	1	1
2,1,1	-1	-1	-1	-1	-1	-1	1	-1	1	1	1	-1	-1	1	1
1,2,1	-1	-1	-1	1	-1	-1	1	1	-1	1	1	1	-1	-1	1
0,3,1	-1	-1	-1	1	1	-1	1	1	1	-1	1	1	1	-1	-1
4,0,0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	1	1	1
3,1,0	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	1	-1	1	1	1
2,2,0	1	-1	-1	-1	-1	-1	1	-1	-1	-1	1	1	-1	1	1
1,3,0	-1	-1	-1	-1	-1	-1	1	1	-1	-1	1	1	1	-1	1
0,4,0	-1	-1	-1	-1	-1	-1	1	1	1	-1	1	1	1	1	-1

The matrix does not have a pure Nash equilibrium. Ann's moves (0, 0, 4) and (0, 1, 3) are weakly dominated. For Beth, the list of weakly dominated moves is longer, namely (2, 0, 2), (0, 2, 2), (3, 0, 1), (0, 3, 1), (4, 0, 0), (3, 1, 0), (1, 3, 0), and (0, 4, 0). After eliminating them, we get the matrix

	0,0,4	1,0,3	0,1,3	1,1,2	2,1,1	1,2,1	2,2,0
1,0,3	-1	-1	1	-1	-1	-1	-1
2,0,2	-1	-1	-1	1	-1	1	-1
1,1,2	-1	-1	-1	-1	-1	-1	-1
0,2,2	-1	1	-1	1	1	-1	-1
3,0,1	-1	-1	-1	-1	1	1	1
2,1,1	-1	-1	-1	-1	-1	1	-1
1,2,1	-1	-1	-1	-1	1	-1	-1
0,3,1	-1	-1	-1	1	1	1	1
4,0,0	-1	-1	-1	-1	-1	-1	1
3,1,0	-1	-1	-1	-1	-1	-1	1
2,2,0	1	-1	-1	-1	-1	-1	-1
1,3,0	-1	-1	-1	-1	1	-1	1
0,4,0	-1	-1	-1	-1	1	1	1

Now Ann's moves (2, 0, 2), (1, 1, 2), (3, 0, 1), (2, 1, 1), (1, 2, 1), (4, 0, 0), (3, 1, 0), (1, 3, 0), and (0, 4, 0) and Beth's moves (1, 1, 2) and (2, 1, 1) are weakly dominated. We eliminate them and get:

	0,0,4	1,0,3	0,1,3	1,2,1	2,2,0
1,0,3	-1	-1	1	-1	-1
0,2,2	-1	1	-1	-1	-1
0,3,1	-1	-1	-1	1	1
2,2,0	1	-1	-1	-1	-1

Since in the reduced submatrix, Beth's moves (1, 2, 1) and (2, 2, 0) have identical payoffs against all of Ann's moves, it suffices to take one of them, say (1, 2, 1). We get the following matrix, which now doesn't

have any weakly dominated moves. It is the result of the IEWD procedure.

	0,0,4	1,0,3	0,1,3	1,2,1
1,0,3	-1	-1	1	-1
0,2,2	-1	1	-1	-1
0,3,1	-1	-1	-1	1
2,2,0	1	-1	-1	-1

What would Ann and Beth choose? Is any of Ann's four options better or even different from the others? The matrix is symmetric in the sense that each of Ann's remaining options beats exactly one of Beth's remaining options, and each of Beth's moves beats exactly three of Ann's remaining options. This may imply that Ann and Beth each choose any of their four remaining options, (1, 0, 3), (0, 0, 2), (0, 3, 1), and (2, 2, 0) for Ann and (0, 0, 4), (1, 0, 3), (0, 1, 3), and (1, 2, 1) for Beth. You may agree that Beth has an advantage in this game. Therefore, having an advantage of 1 in the small districts C and D is better than an advantage of 1 in the large district E. We finish this example in the second part, where we use the tool of mixed strategies.

7.3 The General Model

The general formulation of these games may be as follows:

ELECTION($c, d, e | c_1, d_1, e_1 | a, b$) There are three districts, C, D, and E, and the president is elected by electoral votes. There are c electoral votes from district C, d electoral votes from district D, and e electoral votes from district E. Districts do not split electoral votes, they vote for the candidate having put the most resources into the district, and they abstain in case of a tie. The two presidential candidates, Ann and Beth, simultaneously decide how to allocate Ann's remaining a resources and Beth's remaining b resources over the districts. Presently Ann leads in districts C, D, and E by c_1, d_1 , and e_1 resources. How would Ann and Beth distribute their resources?

7.4 Third Example

ELECTION(7, 9, 13 | -1, -1, 1 | 4, 4) is like **ELECTION** 2, except that district D has 9 votes instead of 8.

We can do the whole procedure once again, will anything change?

To answer this, let me reformulate who wins in the games

ELECTION($c, d, e | c_1, d_1, e_1 | a, b$) after both players have moved simultaneously:

1. If after the additional resources have been allocated all three districts are tied, then the game ends in a tie.
2. If at the end two districts are tied, the third district decides the president.
3. If one district ties, then the player winning the larger district becomes president. If both districts have equal size and are won by different players, then there is an overall tie. (In **ELECTION**(7, 8, 13 | $c_1, d_1, e_1 | a, b$) as in **ELECTION**(7, 9, 13 | $c_1, d_1, e_1 | a, b$) no districts have the same size.)

4. If no district is tied, we add the votes and compare. Since $7 + 8 > 13$, and $7 + 9 > 13$, in $\text{ELECTION}(7, 8, 13|c_1, d_1, e_1|a, b)$ and in $\text{ELECTION}(7, 9, 13|c_1, d_1, e_1|a, b)$ the player winning (at least) two districts wins the presidency.

The sizes of the districts matter only in the third and fourth case, and since $\text{ELECTION}(7, 9, 13|-1, -1, 1|4, 4)$ behaves just like $\text{ELECTION}(7, 8, 13|-1, -1, 1|4, 4)$, both are played exactly the same way. Although their descriptions are different, one could say that the games are the same—we have the same moves and the same payoffs. Mathematicians call such games *isomorphic*.

7.5 The Eight Cases

Let us look again at the four cases. In case 3, what is important is only whether some of c , d , or e are the same. The exact description of case 4 is:

4. Assume $c \leq d \leq e$, and no district is tied. If $c + d > e$, then the player winning (at least) two districts wins the presidency. If $c + d = e$, then the player winning (at least) two districts wins the presidency, except when one wins C and D and the other E, in which case we get an overall tie. If $c + d < e$, then the player winning district E wins the presidency.

Combining the properties on repeated sizes and $c + d$ versus e , we get eight cases.

- 13-13-13: all the districts have the same value,
- 8-8-15: there are two identically sized smaller and one larger district, and the sum of the small district sizes is greater than the size of the largest one,
- 8-8-16: there are two identically sized smaller and one larger district, and the sum of the small district sizes is equal to the size of the largest one,
- 8-8-17: there are two identically sized smaller and one larger district, and the sum of the small district sizes is smaller than the size of the largest one,
- 8-13-13: there are two identically sized larger and one smaller district,
- 7-8-13: all sizes are different, and the sum of the two smaller sizes is larger than the size of the largest one,
- 5-8-13: all sizes are different, and the sum of the two smaller sizes is equal to the size of the largest one,
- 4-8-13: all sizes are different, and the sum of the two smaller sizes is smaller than the size of the largest one.

Our examples belong to the group “7-8-13”.

7.6 Voting Power Indices (optional)

Both players will treat district D the same, no matter whether it has 8 or 9 (or 10 or 11 or 12) votes, provided district C has 7 votes and district E has 13. The number of votes is not proportional to the importance of the district. There is some connection to “voting power indices”, but not an immediate one. The indices apply to situations where voters have different numbers of votes. However, the power of a voter is usually not proportional to the number of votes. Take the example of three voters with 4, 8, and 13 votes: the two voters with the small numbers of votes have no power. But two important assumptions not valid for our model are that no voter abstains—voters can vote only for or against a proposal, and if they can abstain, the abstentions count as votes against—and that no tie is possible, the proposal is accepted if it gets more than half of the votes, and otherwise rejected.

The **Banzhaf-Penrose Index** is calculated as follows: A set of voters is called a **winning coalition** if they have enough votes to pass a proposal. In the example where A, B, C , and D have 4,3,2, and 1 votes, and where six votes are needed to accept a proposal, winning coalitions are $\{A, B\}$, $\{A, C\}$, $\{B, C, D\}$, and every set containing any of them, namely $\{A, B, C\}$, $\{A, B, D\}$, $\{A, C, D\}$, and $\{A, B, C, D\}$. A **swing voter** is a member of a winning coalition whose removal makes the group non-winning. We list the seven winning coalitions again, this time underlining swing voters. We get $\{\underline{A}, \underline{B}\}$, $\{\underline{A}, \underline{C}\}$, $\{\underline{B}, \underline{C}, \underline{D}\}$, $\{\underline{A}, \underline{B}, \underline{C}\}$, $\{\underline{A}, \underline{B}, \underline{D}\}$, $\{\underline{A}, \underline{C}, \underline{D}\}$, and $\{A, B, C, D\}$. Now the A occurs 5 times among the 12 swing voters, B and C occur three times, and D occurs only once. Therefore the Banzhaf-Penrose Indices of A, B, C , and D are $5/12, 3/12, 3/12$, and $1/12$.

In the **Shapley-Shubik Index** we look at all ordered lists, as A, B, C, D or C, D, A, B where each voter occurs exactly once. For each list we include the voters in the list until we get a winning coalition. The last voter added, changing a losing coalition into a winning one, is called “pivotal” for that sequence. Let us list all 24 sequences, underlining the pivotal element in each case. We get

$A, \underline{B}, C, D$	$B, \underline{A}, C, D$	$C, \underline{A}, B, D$	$D, A, \underline{B}, C$
$A, \underline{B}, D, C$	$B, \underline{A}, D, C$	$C, \underline{A}, D, B$	$D, A, \underline{C}, B$
$A, \underline{C}, B, D$	$B, C, \underline{A}, D$	$C, B, \underline{A}, D$	$D, B, \underline{A}, C$
$A, \underline{C}, D, B$	$B, C, \underline{D}, A$	$C, B, \underline{D}, A$	$D, B, \underline{C}, A$
$A, D, \underline{B}, C$	$B, D, \underline{A}, C$	$C, D, \underline{A}, B$	$D, C, \underline{A}, B$
$A, D, \underline{C}, B$	$B, D, \underline{C}, A$	$C, D, \underline{B}, A$	$D, C, \underline{B}, A$

The number of sequences where a voter is pivotal divided by the number of all sequences is the voter’s Shapley-Shubik Index. So we get $10/24, 6/24, 6/24$, and $2/24$ for A, B, C , and D . By coincidence, the numbers are identical to the Banzhaf-Penrose indices. This is usually not the case.

Exercises

1. Use Excel sheets [Election.xlsx](#) and [SIMzero.xlsx](#) to find the IEWD matrix for the game $\text{ELECTION}(8, 8, 15|1, -1, 0|3, 3)$.
2. Use Excel sheets [Election.xlsx](#) and [SIMzero.xlsx](#) to analyze the game $\text{ELECTION}(5, 8, 13|0, 0, 0|3, 3)$.
3. Use Excel sheets [Election.xlsx](#) and [SIMzero.xlsx](#) to find the IEWD matrix for the game $\text{ELECTION}(5, 8, 13|0, 0, 0|4, 3)$.
4. Find the Banzhaf-Penrose Index and the Shapley-Shubik Index of four parties in a parliament, where party A has 44 seats, party B has 33, party C has 15, and party D has 8.
5. Find the Banzhaf-Penrose Index and the Shapley-Shubik Index of four parties in a parliament, where party A has 45 seats, party B has 30, party C has 21, and party D has 4.

CHAPTER 8

Theory 3: Sequential Games I: Perfect Information and no Randomness

“Life can only be understood backwards, but it must be lived forwards.”

— Søren Kierkegaard

Example 1 NIM(6) Six stones lie on the board. Black and White alternate to remove either one or two stones from the board, beginning with White. Whoever first faces an empty board when having to move loses. The winner gets \$1, the loser loses \$1. What are the best strategies for the players?

Student Activity Try your luck in applet [AppletNim7](#) against a friend (hit the “Start new with 6” button before you start). Or play the game against the computer in [AppletNim7c](#). Play ten rounds where you start with seven stones. Then play ten rounds where you start with nine stones. Then play ten rounds where you start with eight stones. Discuss your observations.

In this chapter we look at a class of simple games, namely sequential games. They are games where the players move one after another. Among them we concentrate on games of perfect information, where players know all previous decisions when they move. Randomness will be included after the next chapter. We will learn a little terminology, see how to display the games either as a game tree or a game digraph, and how to analyze them using “backward induction” provided the game is finite. We conclude by discussing whether the solution found by backward induction would be what real players would play, by discussing another approach for sequential games, by discussing the special roles two-person zero-sum games play here, and by discussing briefly the well-known sequential games chess, checkers, and tic-tac-toe.

8.1 Extensive Form: Game Tree and Game Digraph

A sequential game is a game where the players move one after another. A **position** is a situation where a player has to make a **move**, a decision, choosing one of a number of possible moves. In a sequential game a position is linked to just one of the players. Later, when we allow randomness, positions may be situations where a random experiment is performed, so later positions may also be random (or belong to a “random player”). In addition to the positions belonging to some player, there are positions where the game ends, called **end positions**. They are the outcomes; nobody moves after the game has ended. At the end positions, the payoff for each player is given. Every game starts with a **start position**.

The relevant information known to the player about to move is also part of the position. Usually in a sequential game of perfect information the information is the sequence of the decisions of the other players in

all previous moves. Different sequences of decisions imply that we have different positions, except in cases where the difference is not relevant for the future. We will explain this more thoroughly when we explain the difference between game trees and general game digraphs.

The game changes its position by moves of players. At each non-end position, the player linked to the position chooses a move. The possible moves are also known to everybody. Each leads to another position.

Example 2 SENATE RACE II An incumbent senator (from a rightist party) runs against a challenger (from a leftist party). They first choose a political platform, leftist or rightist, where the senator has to move first. If both choose the same platform, the incumbent wins, otherwise the challenger wins. Assume that the value of winning is 10, the value of compromising their political views (by choosing a platform not consistent with them) is -5 , and the payoff is the sum of these values [Kn.d.].

There are four outcomes: If both choose “leftist”, the incumbent senator wins 10 but loses 5 for compromising her views, so the payoff is 5. The challenger doesn’t win anything and doesn’t lose anything in this case, so the payoff is 0. In the same way, if both choose “rightist”, then the senator gets a payoff of 10 and the challenger gets -5 . If the senator chooses “leftist” and the challenger “rightist”, then the senator “wins” -5 (loses 5) and the challenger gets a payoff of 5. Finally if the senator chooses “rightist” and the challenger “leftist”, then the (now former) senator gets a payoff of 0 and the challenger gets 10. Before the four end positions, there is the start position where the incumbent senator has to decide between “leftist” and “rightist”, and two positions where the challenger has to choose. One position is where the senator has chosen “rightist”, and the other where the senator has chosen “leftist”.




Figure 8.1. The extensive form of SENATE RACE II

Sequential games can be described graphically: For each position a small circle is drawn on the paper, and the possible transitions between the positions are expressed by arrows between the circles. Often, instead of arrows we use straight lines or curves, from left to right or from up to down, with the convention that the direction is rightwards (in the first case) or downwards (in the second). We usually use lines from left to right. The small circles are **vertices** or nodes, and the arrows are **arcs**. The vertices and arcs are usually labeled. At each non-end vertex, there is a label indicating to whom the vertex belongs. The outgoing arcs of a non-end position correspond to the moves a player can make at the position. They are labeled by the move names. End positions are vertices with no outgoing arcs, and at them the payoffs for each player are noted. The whole (directed) graph with all labels is called the **game digraph** or **extensive form** of the sequential game. We will later extend the definition, when we take into account randomness and non-perfect information.

The extensive form of SENATE RACE II is in Figure 8.1.

Example 1 The extensive form for the NIM(6) game is in Figure 8.2. If, following chess terminology, the first mover is called “White” and the other one “Black”, then White’s positions, moves, and payoffs




Figure 8.2. The game tree of NIM(6)




Figure 8.3. The game digraph of NIM(6)

are colored red, and Black’s positions, moves, and payoffs blue. Every position has a label indicating how many stones are still on the board at that position, and whose move it is. For example, “4W” is the position with four stones on the board and White to move.

The extensive forms in the previous two examples are **game trees**. We’ll avoid a formal definition and say that a game tree looks like a tree, rotated by 90 or 180 degrees. Trees arise if we consider a position to be the whole sequence of previous decisions.

There is redundancy in the example. Why do we have two “3W” positions? Granted, they have different histories, one resulting from position “5B” with Black taking two stones, and the other from position “4B” with Black taking one stone. But that is not relevant for the future, as can be seen in the game tree: the subsequent subtrees are identical. So White should play the same move in both positions.

If we identify corresponding positions in the game tree of the Nim(6) example, we get the extensive form of Figure 8.3:

It should have become clear that games may have different descriptions by extensive forms. In the literature, game trees are mostly used to describe extensive forms. Using game digraphs often reduces the number of positions.

8.2 Analyzing the Game: Backward Induction

8.2.1 Finite Games

Next we discuss a simple and powerful method for analyzing sequential games. It works only for finite games, those with finitely many options that end after a finite number of moves. A sequential game is **finite** if it has a game tree with finitely many vertices. Having a game digraph with finitely many vertices is not sufficient, as can be seen in the next example:

Example 3 2 FINGERS Two players move alternately. A player moves by raising one or two fingers. A player loses when raising the same number of fingers as the other player in the previous move. Then the payoffs are -1 for the loser and 1 for the winner. If the player shows a different number, the game continues.




Figure 8.4. Infinite game tree of 2 FINGERS

How do you play this simple zero-sum game? Who will win? Obviously nobody will lose, since losing can easily be avoided. So if nobody loses, nobody wins. The two players will continue playing forever. The game tree goes on and on, indicated by the dashed line in Figure 8.4, and is therefore infinite.

We can also use a game digraph. A non-end position is uniquely determined by who is about to move, White or Black, and how many fingers were last shown. We have four positions, labeled as W1, W2, B1, and B2. There is also the start position with White to move, where no fingers have been shown yet, and two end positions, one where White wins and one where Black wins. See Figure 8.5 for the game digraph.




Figure 8.5. game digraph of 2 FINGERS

Why is the game not finite, although it has a finite digraph? The reason is that the game digraph is **cyclic**. It is possible to follow arcs from a vertex x and return to it, as in W1, B2, W1 in the above example. Games with a cyclic game digraph always have an infinite game tree. A sequential game is finite if it has an **acyclic** (not cyclic) finite game digraph.

There are theoretical reasons why we exclude these infinite games, but there are more practical reasons: We need to have a guarantee that a game eventually ends. Even chess would not be finite, were it not for the 50-move rule that says that a chess play is draw if within the last 50 moves no pawn has moved and no piece has been taken.

8.2.2 The Procedure

Players want to know how to play best. In a sequential game that means that for every position that is not an end position the player who has to move wants to know which option to choose. A list of recommendations for all the player’s positions—even for those positions that would never occur—is called a **pure strategy** for that player. In this section we will present a procedure that generates pure strategies for all players. It attaches numbers to vertices: the likely payoffs for the players.

Let us explain the method with another example, this time a non-zero sum game with three players:

Example 4 SEQUENTIAL LEGISLATORS VOTE: Three legislators vote whether they allow themselves a salary raise of \$2000 per year. Since voters are observing the vote, a legislator would estimate the loss of face by having to vote for a raise as \$1000 per year. A has to vote first, then B, then C, and all votes are open. (This is a variant of a game described in [K2007].)

The game tree is in Figure 8.6. The behavior of legislator C is easiest to explain first, since when C has to move, A and B have already moved. So C doesn’t need to anticipate their strategies. C may face four situations, corresponding to the four vertices to the right of the tree. In the first, when A and B have already voted for a raise, C can save face by voting against the raise but still get the benefit of it. The same is true if A and B both rejected the idea of a raise and it is decided already. It is different if one voted for and one against a raise. Then C’s vote counts. Since the money is more important than saving face (for our fictitious legislators only!), C would vote for a raise.




Figure 8.6. Game of SEQUENTIAL LEGISLATORS’ VOTE

Everybody, including B, can do this analysis. When B has to decide, she can anticipate C’s reactions. If A voted for a raise, B knows that C will vote for a raise if B votes against a raise, and conversely. So why not give C the burden of having to vote for the raise? Obviously B will vote against a raise in this situation. If, on the other hand, A voted against a raise, then B has to vote for a raise to keep things open.

Since A knows all this too, A will vote against a raise, B will vote for it, and C will vote for it as well. A has the best role in this game. Going first is sometimes useful!

This method of attaching a recommendation and likely payoffs for players to the vertices of the extensive form, starting at the vertices late in the game and moving backwards is called “backward induction”. We will explain it more formally below.

When a player accepts the decisions found by backward induction, he or she must assume that the other players will also stick to these backward induction recommendations. And why wouldn’t they? One reason could be that they are not able to analyze the game, since it may be too complicated. Or a player can analyze the game, but doubts that all others can.

In NIM(6) we can immediately assign likely payoffs to the vertices 1W and 1B. In both cases, the player to move has just one option. In the 1W case, White will move, take the remaining stone, and arrive at 0B, which is a loss for Black, therefore a win for White. Therefore White will win when facing 1W, and the expected payoffs are 1 for White and -1 for Black. In the same way, the likely payoffs at 1B are -1 for White and 1 for Black.

Having assigned values to four vertices, we consider the vertices 2W and 2B, all of whose successors already have values. In the 2W position, White can proceed to 0B and win, and will do so, therefore the values there are 1 and -1 for White and Black. Similarly, at the position 2B the likely payoffs are -1 and 1 for White and Black.

Next we look at the positions 3W and 3B. From 3W, White can move to position 2B or 1B by taking one or two stones. Both are unfavorable for White, with a likely payoff of -1. Therefore it doesn’t matter what White chooses, and the likely payoffs at vertex 3W are -1 and 1 for White and Black. The position 3B is analyzed in the same way.

From position 4W positions 3B and 2B can be reached. Because the first is favorable for White and the second is favorable for Black, White will choose the first option. This means that the likely payoffs of position 3B are copied to position 4W. In the same way, Black will move to position 3W from position 4B, and the likely payoffs of position 3W are copied to position 4W.

We proceed in this way, and assign likely payoffs to all vertices including the start vertex. They are shown in Figure 8.7. At the start, White expects to lose and Black expects to win.




Figure 8.7. Backward induction analysis of NIM(6)

Procedure for Backward Induction *Your goal is to have likely payoff values for each player assigned at every vertex. As long as not all vertices have values attached, do the following:*

- *Find a vertex with no values attached, but whose successors have values attached. Such a vertex can always be found (why?). Call it V .*
- *At vertex V , one player, say X has to move. Identify the successor vertex of V for which the value for player X is highest. This is the vertex to which player X wants to move (since the values will turn into payoffs eventually). Let's call this other vertex W .*
- *Copy the values of vertex W to vertex V . The backward induction strategy recommends that in position V player X will move in such a way that position W is obtained.*

There is one problem in this procedure. We have not determined how to handle ties, if two or more successor vertices of a vertex have the same likely payoff for the player to move. For this, see Subsection 8.3.4.

8.2.3 Zermelo's Theorem

In the literature, the following theorem is often named after Ernst Zermelo, who published a paper on perfect-information 2-player zero-sum games [Z1913] in 1913. However, Zermelo didn't apply backward induction, [SW2001], and was mostly interested in giving a bound on the number of moves needed to get a win in a winning position. Still, we will follow the convention and call the following Zermelo's Theorem:

Theorem (Zermelo) [Z1913] *Every finite perfect-information sequential game without random moves can be analyzed using backward induction. The payoff values of the start position are what the players should expect when playing the game, provided all play rationally.*

By construction, the outcome generated by backward induction has the following property, which has some connection to the Nash equilibrium definition:

Fact *In the backward induction strategies, if a player deviates from his or her backward induction strategy, then the player's payoff will not increase.*

Historical Remark Ernst Zermelo (1871–1953) was a German mathematician. He was Professor in Göttingen from 1905 to 1910 and in Zürich until 1916, when he had to retire because of health problems. He is known best for his work in the foundations of set theory, invented by Georg Cantor about 30 years earlier. However, his paper [Z1913] was his only contribution to game theory.

8.3 Additional Topics

8.3.1 Reality Check

Although the backward induction solution seems to be the best possible for players, there are sequential games where players tend to play differently:

Example 5 CENTIPEDE GAME In this game, two players alternately face two stacks of money. To make a move, a player has the choice either to pass, in which case both stacks grow slightly and the other player now must make a move facing slightly larger stacks, or to take the larger stack, in which case the other player gets the smaller one and the game ends. If it didn't end before, the game ends after a fixed number of rounds, in which case both players share the accumulated money evenly.

The example in Figure 8.8 shows the version where initially the two stacks contain 2 and 0 units, increase by 1 unit in each round, and the game must end after at most six rounds.




Figure 8.8. game tree of CENTIPEDE GAME

The game was introduced by Rosenthal [R1981] in a 100 rounds version. If the increment amount is smaller than the difference between the two stacks, it has a unique backward induction solution: Player 1 would take the money in the first round and end the game. However, players can increase their payoff if they both wait. When McKelvey and T. Palfrey made experiments with this game [MP1992], they observed that players tend to play some rounds until one of them cannot resist temptation any longer and takes the larger stack.

Here is another classical example:

Example 6 ULTIMATUM GAME There is a fixed number of dollar bills for both players, say 5. Ann makes an offer how to share it, which Beth can either accept or reject. If she accepts, the bills are divided as agreed upon; if she rejects, nobody gets anything.

Figure 8.9 shows the extensive form of the game where the payoffs are the money obtained. In the backward induction solution Ann would offer only one bill to Beth, or maybe none. We will discuss this problem with payoff ties in Subsection 3.4. This is another example where people usually deviate from theory. Experiments show that people like to be fair. Ann usually offers an amount below but close to one half of the total amount, but if Ann offers too little, Beth usually gets furious and refuses the offer, leaving both with nothing.

Modeling Note Modeling Altruism using Utilities Working through the material so far, you may be annoyed with these rational players. They may appear to you rather cold, or inhuman. Granted, companies, aiming to maximize profits can be modeled by rational players, and granted that in parlor games, in particular those that are zero-sum, players may consider only their own payoff. But don't we as humans, in real-world situations care about each other? Doesn't that mean that we are not only looking at our own payoffs, and are not indifferent about the payoffs of others?




Figure 8.9. Game tree of ULTIMATUM GAME

Things may become clearer if we separate money paid from satisfaction received. Suppose a player’s payoff is his or her satisfaction, instead of money received. The payoff depends on all features of the outcome, including the money given to all players. The dependence is often called the utility function. It is usually different for different players. Only for a very selfish player is payoff or satisfaction the same as the money he or she obtains. The payoff of a totally altruistic player could be the sum of the money won by all the players, and there are possibilities between the extremes. In this way, altruism, empathy, as well as antipathy, can be modeled. In most cases the utility function is not linear. Winning twice as much money doesn’t make you twice as happy. See Chapter 18 for more about nonlinear utility functions.

Let us try to model dissatisfaction with unfairness as follows. Let the payoff of each player, the satisfaction, be the money obtained, reduced by an unfairness penalty for the player obtaining more. Assume this unfairness penalty is two thirds of the difference between both values. That means that the money exceeding the amount the other gets is still valuable, but only one third of this surplus money is satisfying for the player. Payoffs in this modified Ultimatum Game can be seen in Figure 8.10. The backward induction solution has Ann offering two bills to Beth, which she accepts.

8.3.2 Playing it Safe—Guaranteed Payoffs

A player will get the payoff predicted by backward induction analysis only if all players play rationally, i.e., if they make the backward induction moves. Thus this likely payoff is not guaranteed: if the other players deviate from backward induction moves, a player may get less. Therefore using backward induction moves may be risky. If you want to play it safe, maybe by assuming that all the other players want to hurt you, as in the maximin paradigm discussed in Chapter 2, you may want to play moves different from those prescribed by the backward induction analysis. See the example given by Figure 8.11.




Figure 8.11. Another game tree

In this game Ann starts, there are two positions where Beth is supposed to move, and there are four end positions. The backward induction analysis goes as follows: In Beth’s first position, Beth will choose move B_1 , and in her second position she will choose move B_3 . Therefore the likely payoffs at these two positions are 1, 10 and 2, 9. Therefore Ann will choose move A_2 , since this will give her a payoff of 2. However, if Ann doubts Beth’s rationality, and fears that in Beth’s second position she may choose B_4 , then Ann may choose A_1 , since this guarantees her a payoff of at least 1.

Sequential games with perfect information have **guaranteed payoffs** for all players. These are also called **security levels**. They can be found for every player separately, doing a variant of backward induction, which will be called **security level analysis**. Let’s call the player in question Ann. During this analysis we also find what moves Ann would choose in each of her positions if she would fear that everybody is playing just to hurt her. The resulting strategy is called her **security strategy**. Since this takes only Ann’s payoffs into account, we can disregard other’s payoffs. We attach a value, the guaranteed payoff for Ann, to every position, starting again at the end positions. We assign values to vertices whose successors already have values, guaranteed payoffs for Ann, attached. If Ann does not move at that vertex, then we expect the worst and choose the lowest of the values of the successors as the value for the vertex. We cannot guarantee more in that case. If Ann moves at that position, she can move to the successor with the highest guaranteed payoff, so we assign the largest of the values of the successors to the vertex.

In the previous example, Ann’s guaranteed payoff at Beth’s first (upper) position is 1, and Ann’s guaranteed payoff at Beth’s second position is 0. Therefore Ann’s guaranteed payoff at her start position is 1, achieved by move A_1 . This strategy differs from the backward induction strategy.

Consider the game described in Figure 8.12:

- Backward induction analysis indicates a likely payoff of 10 for Ann and of 6 for Beth. Ann would choose A_1, A_4, A_7, A_{10} , and A_{11} , and Beth B_2, B_4, B_6, B_8, B_9 , and B_{11} .
- Security level analysis for Ann shows that Ann can guarantee a payoff of 5 by choosing move A_2 instead of A_1 and A_9 instead of A_{10} .




Figure 8.12. Still another game tree

- Security level analysis for Beth shows that Beth has a security level of 3. Beth would sometimes move differently than in the backward induction strategy.

8.3.3 Two-person Zero-sum Games

Fact *If a game has only two players and is zero-sum, then the backward induction analysis is the same as the security level analyses. The backward induction strategy is the same as the security level strategies of the two players, and the values obtained by backward induction are the security levels of the players. So assuming rational play of the other player, playing safe is the same as playing rationally.*

The reason for this is simple. When we apply the procedures, we look at a position, say of Beth’s, where all successor positions have values assigned. A move minimizes Ann’s values if it maximizes Beth’s values, if the sum of the values is zero.

This fact gives tremendous significance to backward induction in the two-player zero-sum case. Therefore it seems justified to call the backward induction strategies of the players optimal. If Ann and Beth use the strategies, they will obtain the values attached at the start position as payoffs. If one of them, say Beth, deviates, we already know that Beth will get not more. This is true even without the assumption that Ann and Beth are playing a two-player zero-sum game. But with the additional assumption, if Beth deviates from her backward induction strategy, we know that Ann will not get less, provided Ann sticks to her backward induction strategy.

Historical Remark Most people consider John von Neumann (1903-1957) to be the father of game theory. His first paper on the topic appeared in 1928. Although Emile Borel had published some papers on simultaneous games in the 1920s, they, and Ernst Zermelo’s paper mentioned above, were not as influential as von Neumann’s paper. Von Neumann concentrated from the beginning on 2-player zero-sum games that allow

the most impressive theory. In 1944 von Neumann and the economist Oscar Morgenstern published their fundamental book *The Theory of Games and Economic Behavior* [VNM1944].

Born in Budapest, von Neumann studied and worked in Zürich and Berlin. In the 30s he emigrated to the USA, where he had been offered a professorship at the newly-founded Institute of Advanced Study in Princeton. During World War II he played a leading role in the Manhattan project, which developed the atomic bomb, and he was an influential US government consultant during the Cold War. In the 40s von Neumann played an important role in the development of the first computers. See [D2013] for more information.

8.3.4 Breaking Ties

What happens in the backward induction procedure if there are two or more successor vertices, say W and U , promising both the maximum value for player X to move? Player X would be indifferent whether to move to W or to U . The value for X at V is obviously the value for X at W (which is the same as the value for X at U), but the decision of where to move would influence the payoffs of the other players. It is natural to assume that X will move to W or U with equal probability. This is a behavioral strategy—at certain vertices a player randomly chooses from two or more moves. Then we call X a neutral player, and the value for other players at V would be computed as the average (expected value with 50-50 probabilities) of the values for that player at W and U .

In a two-person game, a different approach could be taken. If player X is hostile to her opponent, she would choose a vertex of W and U that has the smaller value for her opponent. If X is friendly she would choose the vertex of W and U that has the larger value for her opponent. She still cares only about her payoffs, only in a tie situation of her own payoff would she consider the other's payoff. Hostile players can be modeled by defining new payoffs for them, formed by subtracting a small fraction (let's say 0.1%) of her opponent's payoffs from her former payoffs, a fraction so small that the ordering of the payoffs is not changed. Then in ties the situations are ordered. Similarly, a friendly player would add a small fraction of the opponent's payoffs to her own former payoffs.

There are more ways of breaking ties, even among more than one player. A player may be friendly to some players and hostile to others. Backward induction analysis must be changed in each of these cases.

8.3.5 Existing Games

Board games are, almost by definition, games of perfect information, since everything is in the open provided cards or dice are not part of the game. Often they are also sequential games, where the players take turns to move. Well-known examples of board games are chess, go, nine men's morris, checkers, parcheesi, and backgammon (which has a very long history). More recent games are monopoly, hex, and tic-tac-toe. Some (for example parcheesi, backgammon, and monopoly) are played with dice or with cards that are revealed during play, and so contain randomness. But chess, go, checkers, tic-tac-toe, hex are two-player zero-sum finite sequential games without randomness and with perfect information. So shouldn't people play the backward-induction strategy? Yes, they should, but in some cases the backward induction analysis has not yet been, or cannot be, performed. Thus the optimal strategy is unknown in many cases.

Tic-tac-toe is the simplest of these games, with the smallest game tree. The game tree has just 5478 positions. It is not too difficult to perform backward induction to see that the expected outcome is a draw. Nine men's morris is more complex, with about $10^{10} = 10,000,000,000$ positions. In 1994, Gasser and Nievergelt [G1996] [GN1994] showed that it will also end with a draw, when played optimally by both players. Connect-four is a game with about $5 \cdot 10^{12}$ positions. It was shown independently by James Dow

Allen and Victor Allis [A1988b], [A1989] in 1988 that White must win, when playing optimally.

Much more complicated is checkers, with about $5 \cdot 10^{20}$ positions, 100,000,000 times as large as the number of positions of Connect-four. Full backward induction seems out of reach—it is not even feasible to store all these positions. Nevertheless, by looking only at a small part of the game tree, a group led by Jonathan Schaeffer proved that checkers will always result in a draw when played correctly [SBBKMLLS2007].

Chess has about 10^{43} positions, a number with about 43 digits; and go about $2 \cdot 10^{170}$. Neither will be solved soon.

8.3.6 Greedy Strategies

How do computers play chess? The complete extensive form is not available—it is too large—so backward induction is not an option. Existing chess programs usually have two ingredients. Starting at a position, they look a few moves ahead, instead of looking back. Secondly, they can attach values to every possible position. The values are supposed to be approximations of the likely payoffs found by backward induction, but they are not obtained by looking further to successors of the position considered. Rather features of the position are considered, and a value for the position is computed from a complicated formula. In chess, material advantage is one of the features these programs consider, but there are also others, like the number of possible moves from a position, the pawn structure, and so on.

A **greedy strategy** decides what to do in a position by looking at successor positions and choosing the one where the evaluation for the player to move is highest. More generally, a **k -Greedy strategy** looks k moves ahead, and chooses the maximin or backward induction choice of this small part of the game digraph.

Exercises

1. Draw the extensive form of the following game:

TWO-ROUND BARGAINING(5, 3) Five dollar bills are to be shared by two players. Ann makes an offer how to share them, which Beth can either accept or reject. If she accepts, the bills are divided as agreed upon. If she rejects, two dollar bills are taken away and only three remain on the board. Then Beth makes an offer how to share them, which Ann can accept or reject. If Ann accepts, the bills are divided accordingly and if she rejects nobody gets anything.

Draw the extensive form of the game, and perform backward induction analysis. Remember that you don't have to use a game tree—you can use a digraph to simplify matters.

2. Analyze the following game:

MATCHING CHAIRS Ann and Beth can select both three chairs out of seven available, two L-chairs and five O-chairs. One L-chair is worth \$300, but if you have a pair of L-chairs, the pair is worth \$800. O-chairs are less valuable: One is worth \$100, a pair is worth \$400, but three are worth \$900. Beginning with Ann, Ann and Beth alternate selecting a chair until each has three.

Draw the extensive form of the game, and perform backward induction analysis.

3. Analyze the following game:

REC THE SQUARE Two players, White and Black, alternately put dominos (where a domino is a rectangle consisting of two adjacent squares) on a 3×3 checkerboard. The player who cannot fit a domino in loses, the other one wins.

You can play the game in the applet [RecTheSquare3](#). Draw the extensive form of the game, and perform backward induction analysis. You can describe the game using the 18 states shown in Figure 8.13, where each symbol subsumes the cases obtained by rotation and reflection.




Figure 8.13. The 18 states for REC THE SQUARE

4. Analyze the following variant of REC THE SQUARE:

CREATING HOLES: White and Black alternately put dominos on a 3×3 checkerboard. Each player has to move until there is no domino space left. Then White gets \$1 for every empty corner field, and Black gets \$1 for every other empty field.

You can play the game in the applet [RecTheSquare3](#). Draw the extensive form of the game, and perform backward induction analysis. You can describe the game using the same 18 states shown in Figure 8.13.

The next two exercise refer to the following game:

TAKE SOME(n, p, q) At the beginning n dollars lie on the board. p is a number between 0 and 1, and q is a positive number. There are five rounds, where, starting with Ann, the two players that have not yet “finished” alternately move. A player who is to move can either take part of the money on the board, or wait. A player can take money only once. A player who takes some money has finished and cannot move anymore. If one player has finished, the other one moves every round. In the first round, the player who is to move can take a fraction p of the money, in the second round $p - 10\%$, in the third round $p - 20\%$, and so on. In rounds where no money is taken, the money on the board increases by $\$q$.

The money increases if the players wait, but the fraction they are allowed to take decreases over time, and moreover the money decreases when a player takes some. The trick is to find the right time to take money.

- 5. Analyze TAKE SOME(8, 60%, 10). Draw the extensive form of the game, and perform backward induction analysis.
- 6. Analyze TAKE SOME(12, 60%, 10). Draw the extensive form of the game, and perform backward induction analysis.
- 7. Find the backward induction solution and perform security level analysis of the game in Figure 8.14.




Figure 8.14. A game tree

- 8. Find the backward induction solution and perform security level analysis of the game in Figure 8.15.




Figure 8.15. A game tree

- 9. Find the backward induction solution and perform security level analysis of the game in Figure 8.16.

The next two exercises refer to the following game:

WHO'S NEXT(n) Three players play a sequential game with n moves. At the beginning Ann holds a dollar bill. In each round the player having the dollar bill has to give it to one of the other two players. After n rounds, after the dollar bill has moved n times, the game is over and the player holding it can keep it.

- 10. Draw the game digraph (rather than a game tree, which would be too large) of WHO'S NEXT(5).
- 11. Draw the game digraph (rather than a game tree, which would be too large) of WHO'S NEXT(6).
- 12. Draw the extensive form, a game digraph rather than a game tree, of the following game. There are four players, A, B, C, and D, sitting in a cyclic order. At the beginning A has a dollar bill. In each round the




Figure 8.16. A game tree

player having the dollar bill has to give it to the player to the right, or to the player sitting opposite, but not to the player sitting to the left. After six rounds, after the dollar bill has moved six times, the game is over and the player having it can keep it.

The next two exercises refer to the following game:

CHOCOLATES IN THE BOX (n) There are four players, A, B, C, and D, sitting in a cyclic order. At the beginning A holds a box containing n pieces of chocolate. In each round the player having the box eats one (just one!) and then has to give the box to the player to the right, or to the player sitting opposite, but not to the player sitting to the left. After n moves, the chocolate is gone. The payoff for each player is the number of pieces of chocolate he or she has eaten.

- 13. Draw the game tree of CHOCOLATES IN THE BOX (4).
- 14. Draw the game tree of CHOCOLATES IN THE BOX (5).
- 15. The next game can be analyzed without drawing the game digraph:

5 PIRATES Five pirates with a strict hierarchy have to decide how to distribute 20 gold coins they have found. They do it according to the following rules. In each round the highest ranking pirate has to propose a distribution of the coins. This proposal is voted on, with the proposer voting too, and deciding in case of a tie. That implies that if there are two pirates, the proposal will always win, and if there are four pirates, the proposer needs only one of the others voting for the proposal. If the proposal wins, everything is settled. If not, the proposer is thrown into the water and drowns, and the next round starts. Assume that the pirates value their life at 40 gold coins, and prefer throwing somebody overboard if everything else is equal (they are playing a hostile variant).

What would you propose if there are 2, 3, 4, or 5 pirates left? You can use backward induction without drawing the game digraph!

16. How does the analysis of the game change if we change the rules such that if the vote is a tie the proposal is rejected.
17. Draw the game tree and analyze the variant of the CENTIPEDE GAME that starts with two stacks of 1 and 2 units, where each stack increases by 2 units in each round, and where the game ends after seven rounds.

Project 12

TAKE SOME Analyze the game $\text{TAKE SOME}(n, p, q)$, described in Exercise 5, in the most general form, with all parameters open. Sketch the extensive form of the game, and explain what will happen if both players play optimally, depending on the parameters.

Project 13

WHO'S NEXT(n) This sequential game, described in Exercise 10, has many payoff-ties. Investigate how small sympathies and antipathies between the players would affect the outcome. What happens if A likes B and C, but B dislikes A and likes C, and C likes A and dislikes B, for instance. There are many cases to consider.

How will the players play? Modify the model if you deem it necessary.

Project 14

LISA'S GAME Find and describe a method that finds for any given sequence the solution of this game, described in the applet [LisaGame](#).

Project 15

2-AUCTION Analyze the game described in the applet [Auction2](#). Start by analyzing the game where only one painting is auctioned, and where both players have different amounts of money (in thousands of dollars) available, since these are the situations the players eventually face when the second painting arrives.

Project 16

3-AUCTION Can the method from the previous project be extended to analyze the game described in [Auction3](#)?

CHAPTER 9

Example: Dividing A Few Items I

Prerequisites: Chapters 1 (and maybe 8).

A few items, five or six, are distributed to Ann and Beth by letting them choose one by one. If they alternate choosing, beginning with Ann, we call the games **ABABA** and **ABABAB**, depending on the number of items. For different choice orders we get different games. For instance, **ABBABA** is a game with six items where Ann chooses one item, Beth chooses two, Ann chooses one, Beth one, and Ann takes the remaining one. In the same way the games **ABBAAB** could be defined for six items, and the games **ABBAB**, **ABBAA**, and **ABABB** for five items.

We assume that the items may have different values to Ann and Beth, that the players know how they both value each item, and that the game's total value to a player is the sum of the values of the items she got. Both players want to get as much value as possible.

Let us label the items as item C , D , $\dots$ and let $a(C)$ and $b(C)$ denote the values of item C for Ann and Beth, and so on.

9.1 Greedy Strategy

Isn't the way how to play these games obvious? Wouldn't each player choose the most valuable remaining item(s) whenever she has to choose? Let's call this strategy the greedy one. Instant gratification now!

Student Activity Try the game yourself, with a friend, in applet [ABABAB](#) for six items. Try in particular the "Example 1" data and the "Example 2" data. Play the same game repeatedly with the same values, once where both of you play as well as you can, once where both play greedily, once where Ann plays greedily and Beth gets and obeys the hints given there, once where Ann gets and obeys the hints and Beth plays greedily, and once where both get and obey the hints. Do this for all three games and answer the questions:

- Is obeying the hints better or worse than the greedy strategy? Always?
- Which of the games seems to favor Ann, which favors Beth?

There are situations where the greedy strategy is not best, and rational players would avoid it. For instance, assume **ABABAB** is played, and three items are left with values of $a(C) = 1$, $a(D) = 2$, $a(E) = 3$ for Ann and $b(C) = 3$, $b(D) = 1$, $b(E) = 2$ for Beth. If Beth plays greedily, she takes item C , Ann selects item E , and Beth gets the remaining item D , therefore Beth will add a value of 4 to the value of her first item. However, knowing that Ann will not take item C , which is of little value to Ann, Beth could delay choosing it and instead select her second choice E first. Then Ann would choose item D (remember that Ann, as a

rational player, is not looking for revenge, but aims to maximize her own payoff), and Beth gets item C , adding a value of 5 to the value of her first item.

Ann's and Beth's last move will always be greedy: they just choose the most valuable item.

9.2 Backward Induction

We analyze the sequential game with perfect information by applying backward induction to the game tree or game digraph.

Let us first discuss the number of moves. If a player has two consecutive choices, as Beth in ABBABA, then they could be combined into one move. Taking the last remaining item is not considered a move, since there is no choice. Thus ABABAB is a five-move game, ABBABA a four-move game, and ABBAAB a three-move game.

9.2.1 Game Tree

A position is essentially the total choice history so far. However, we do not care about the order of selections in double-moves.

- In ABABAB, Ann has six options for her first move. Therefore there are six second-round positions for Beth. In each case, Beth has five options. This leads to $6 \cdot 5 = 20$ third-round positions for Ann, abbreviated as pairs of Ann's choice and Beth's choice, as $(C, D), (C, E), \dots, (C, H), (D, C), (D, E), \dots, (H, G)$. Ann has four options in each third-round position, Beth has three options in round 4, and Ann two options in round 5. Consequently, there is one first round position, 6 positions in round 2, $6 \cdot 5$ positions in round 3, $6 \cdot 5 \cdot 4$ positions in round 4, $6 \cdot 5 \cdot 4 \cdot 3$ positions in round 5, and $6 \cdot 5 \cdot 4 \cdot 3 \cdot 2$ end positions. The total number of positions in the game tree is $1 + 6 + 6 \cdot 5 + 6 \cdot 5 \cdot 4 + 6 \cdot 5 \cdot 4 \cdot 3 + 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 = 1237$.
- In ABBABA, Ann has six options for her first move. Then Beth chooses two items out of the remaining five, and has 10 ways to do it. The rest is as in the previous example, so we get $1 + 6 + 6 \cdot 10 + 6 \cdot 10 \cdot 3 + 6 \cdot 10 \cdot 3 \cdot 2 = 607$ positions in the game tree.
- In example ABBAAB, Ann has six options, then Beth has 10 options, and Ann has 3 options after that—choosing two items out of the remaining three. Therefore the game tree has $1 + 6 + 6 \cdot 10 + 6 \cdot 10 \cdot 3 = 247$ positions.

9.2.2 Game Digraph

Does it matter in round 4 whether Beth faces position (C, D, E) or (E, D, C) ? Does it matter whether Ann selected item C first and item E later or the other way around? Obviously not, so we might want to identify these positions.

What we need is only the distribution of the items so far. We list the items Ann got so far (arbitrarily ordered), and separate them from the list of items Beth got so far by the symbol “|”, as for instance $(C, E|D)$ for a digraph position. In round 2 of ABABAB, there are 6 digraph positions. In round 3, we still have $6 \cdot 5$ digraph positions. But in round 4, there are 15 ways of assigning two items to Ann, and 4 ways of having assigned one of the remaining items to Beth. Thus there are $15 \cdot 5 = 75$ digraph positions at round 4. In the same way, in round 5 there are $15 \cdot 6 = 90$ digraph positions, and there are 20 end digraph positions having assigned 3 items to Ann and 3 items to Beth.

- In ABABAB, we get $1 + 6 + 30 + 60 + 90 + 20 = 207$ digraph positions in the game digraph, compared to 1237 in the game tree.
- In the same way, the number of digraph positions in the game digraph of the game ABBABA is $1 + 6 + 60 + 90 + 20 = 177$.

- Finally, the number of digraph positions in the game digraph of ABBAAB is $1 + 6 + 60 + 20 = 87$.

Once these positions have been generated, drawing the arcs is not difficult. For instance, in ABBAAB at digraph position $(C|F, G)$, Ann is supposed to select two items. Then there is an arc from $(C|F, G)$ to $(C, D, E|F, G)$ but not to $(C, D, E|F, H)$. Unlike game tree positions, digraph positions usually have more than one predecessor. For instance, there are arcs from $(C|F, G)$ to $(C, D, E|F, G)$, from $(D|F, G)$ to $(C, D, E|F, G)$, and from $(E|F, G)$ to $(C, D, E|F, G)$, so all three digraph positions $(C|F, G)$, $(D|F, G)$, and $(E|F, G)$ are predecessors of $(C, D, E|F, G)$. The extensive form is no longer a tree but a digraph.

In both 9.2.1 and in 9.2.2 we can perform backward induction to analyze the game. The shape of the game tree or game digraph is the same, no matter what values the players assign to the items, but the analysis of course depends on the values.

9.2.3 Example: Game Digraph for ABBAB

The game digraphs for the 6-item games are too large to be displayed, so let us illustrate the concept with a 5-item game ABBAB. You can play the 5-item games in the applet [ABABA](#). ABBAB is a 3-round game, and there is one digraph position for round 1, 5 digraph positions for round 2, $5 \cdot 6 = 30$ digraph positions for round 3, and 10 end digraph positions. The game digraph and backward induction for Ann's values $a(C) = 40$, $a(D) = 25$, $a(E) = 15$, $a(F) = 13$, $a(G) = 7$ and Beth's values $b(C) = 7$, $b(D) = 40$, $b(E) = 15$, $b(F) = 10$, $b(G) = 28$ are shown in Figure 9.1. This is the example you get in the applet by pressing the "Example 1" button. We suppress Beth's meaningless last move and we omit commas—position $(DG|CEF)$ should read $(D, G|C, E, F)$.

The result of the backward induction analysis is shown by the values at the vertices, and by the bold lines indicating occurring moves. Ann will not play greedily, but will take item D first. Then Beth will select either items E and F , items E and G , or items F and G . Then Ann will select item C and Beth will take the remaining item (G, F , or E).

Optimal play results in payoffs of 65 and 53 for Ann and Beth, but greedy play would produce payoffs of 53 and 83. One might argue that playing greedily is better for the society as a whole. This will hardly convince Ann to play greedily.


9.3 An Abbreviated Analysis

Even the game digraph of ABABAB with 207 vertices is too large to be drawn here. In this section we will change the notion of a position again. With this changed definition, there are only 64 positions left, reducing complexity.

Digraph positions $(C, D|E)$ and $(C, E|D)$ lead to different outcomes with different payoffs. Still, when Beth faces any of these digraph positions, she will choose the same item, since relevant to her decision is only what items remain. In both cases these are items F, G, H . Only these items determine what Beth (and Ann) can add to the values already accumulated. So instead of assigning payoffs for Ann or Beth to game digraph positions $(C, D|E)$ or $(C, E|D)$, we could assign *gain values* to *set positions*. **Set positions** indicate the sets of items left, listed in any order. The **gain value** for Ann or Beth of a set position $\{F, G, H\}$ would be what can be added to the value Ann and Beth have accumulated, provided they play optimally from that position.

As in the previous example, the set positions are the vertices of a digraph, and every such position except the end positions belongs to one of the players. So we get the **gain digraph** of the game. The one for ABABA is displayed in Figure 9.2 (with commas between the numbers omitted).

We find the gain values attached to set positions using backward induction: The start is at the end vertex, with no item left, where there is no gain for Ann or Beth. The gain values for the five set positions $\{C\}$, $\{D\}$, $\{E\}$, $\{F\}$, or $\{G\}$, where Ann moves, are obvious. Ann gains her value of the corresponding item, and Beth gains nothing. The gain values are obvious for the next level with two remaining values. Beth takes






Figure 9.2. Abbreviated gain digraph for ABABA with Ann’s values 40, 25, 15, 13, 7 and Beth’s values 7, 40, 15, 10, 28 for items $C, D, E, F,$ and G

the item worth more to her, Ann gets the other one, and the values to them are their gain values. In general, the reasoning is as follows: Assume we have a set position, like Ann’s $\{C,D,E\}$, and assume that for all successor vertices both gain values have been found. The gain for Ann of an option (taking an item) consists of the immediate gain of the value of that item plus the gain value of the resulting set position. For Ann’s set position $\{C,D,E\}$, option 1 has a gain of $40 + 15$ (40 is the value of item C for Ann, and 15 is Ann’s gain value in set position $\{D,E\}$), option 2 has a gain of $25 + 40$, and option 3 has a gain of $15 + 40$. Ann chooses the option with the largest gain for her, in this case option 2, and the gain value for Ann is $25 + 40 = 65$. The gain value for Beth in $\{C,D,E\}$ is Beth’s gain value for the set position that will occur next, which is set position $\{C,E\}$, so Beth’s gain value at $\{C,D,E\}$ is 15. Proceeding in this way back to the start vertex $\{C,D,E,F,G\}$, we determine gain values there of 78 for Ann and 43 for Beth. The gain values of the positions are displayed in Figure 9.2.

The Excel sheet [ABABAB.xlsx](#) can be used to analyze the games ABABAB, ABBABA, and ABBAAB for any values. They use the method discussed in this section.

9.3.1 Why it Matters: Complexity (optional)

The advantage of the abbreviated approach may not be apparent yet. We can analyze every sequential game without it, and the abbreviated approach applies only to some special sequential games, the games discussed

in this chapter among them. What happens in the strictly alternating game ABABAB... as the number n of items increases? The following table shows the number of tree game positions, digraph positions, and set positions for some values of n .

n	tree positions	digraph positions	set positions
5	206	66	31
6	1237	207	63
7	8660	610	127
10	6,235,301	16,043	1,023
12	823,059,745	137,821	4,095
15	2,246,953,104,076	3,440,637	32,767

9.4 Bottom-Up Analysis

On game trees or, digraphs, or gain digraphs we can find all backward induction solutions. Often there are more than one, since there are many tie situations. For instance, in the game digraph example in Figure 9.1 for ABBAB and $a(C) = 40, a(D) = 25, a(E) = 15, a(F) = 13, a(G) = 7, b(C) = 7, b(D) = 40, b(E) = 15, b(F) = 10, b(G) = 28$, Ann will choose item D , then Beth will choose any two items out of E, F, G , Ann will choose item C , and Beth will take the remaining item. So here we have three backward induction solutions. In the gain digraph example in Figure 9.2 for ABABA and $a(C) = 40, a(D) = 25, a(E) = 15, a(F) = 13, a(G) = 7, b(C) = 7, b(D) = 40, b(E) = 15, b(F) = 10, b(G) = 28$, Ann will choose item D , then Beth will choose item E , then Ann will choose item F or C , then Beth will choose item G , and Ann will take the remaining item, which is either C or F . So we have two backward induction solutions.

It turns out that one of the backward induction solutions is special, since it can be found easily and quickly, without any of these tree or digraph. The method to find it has been described in [KC1971]. It is based on the observation that no player will ever voluntarily choose the item he or she ranks lowest. The player does this only if there is one item left to choose, but that is not a voluntary choice. Furthermore, as we have seen, players playing strategically sometimes delay taking the best available item provided they know that they can take it later. Let's look at the ABABA example in Section 3, where Ann ranks items in the order C, D, E, F, G from most valuable to least valuable, and Beth ranks them in the order D, G, E, F, C . Although Ann values item C most, she knows that Beth will never choose it, so Ann can wait and take it only at the end. If she plans to do so and both players know it, they play essentially the game ABAB with items D, E, F, G , with preference order D, E, F, G for Ann and D, G, E, F for Beth. Since Beth knows that Ann will never take item G , she can wait and take it on the fourth move. Then we get the game ABA with items D, E , and F , and preference orders D, E, F for both Ann and Beth. Therefore Ann selects item F (the lowest for Beth) in move 3, and so on. The items are chosen in order D, E, F, G, C .

The method described in the example was named **bottom-up analysis** in [BT1999]. The choices are determined one by one from last to first. At each step, the player to choose selects the item that has least value for the other player. As in backward induction analysis we proceed backwards, but the analysis is simpler, since in every step the choice can be named. The bottom-up selection depends only on how each player orders the values of the items, not on the values themselves. So even if the values of the items are not cardinal but only ordinal, the method can be applied.

The choices are plausible, but the proof that they form a backward induction solution is complicated and is omitted. See [KC1971].

9.5 Interdependencies between the Items (optional)

Things change if there are interdependencies between the items. An item could become more useful if another item is also owned, as in the case of two shoes. Such positive effects are called synergies. The possession of one item could also lower the worth of another item, as in the case of two similar computer programs.

Game tree and game digraph approaches with backward analysis will still work. The abbreviated gain-method will no longer work, since gain in selecting an item depends on what is already owned.

Exercises

1. a) Use the game digraph in Figure 9.1 to find all backward induction solutions for ABBAB with values of $a(C) = 3$, $a(D) = 4$, $a(E) = 2$, $a(F) = 5$, $a(G) = 1$ and $b(C) = 4$, $b(D) = 3$, $b(E) = 5$, $b(F) = 2$, $b(G) = 1$.
b) Apply bottom-up analysis.
2. a) Use the gain digraph in Figure 9.2 to find all backward induction solutions for ABABA with values of $a(C) = 3$, $a(D) = 4$, $a(E) = 2$, $a(F) = 5$, $a(G) = 1$ and $b(C) = 4$, $b(D) = 3$, $b(E) = 5$, $b(F) = 2$, $b(G) = 1$.
b) Apply bottom-up analysis.
3. Apply bottom-up analysis for the game ABABABAB and the values $a(C) = 7$, $a(D) = 6$, $a(E) = 5$, $a(F) = 2$, $a(G) = 9$, $a(H) = 3$, $a(I) = 8$, $a(J) = 4$ and $b(C) = 9$, $b(D) = 5$, $b(E) = 8$, $b(F) = 4$, $b(G) = 2$, $b(H) = 3$, $b(I) = 7$, $b(J) = 6$.
4. For three players, bottom-up analysis no longer works, but the gain digraph method does. Modify the gain digraph in Figure 9.2 to find all backward induction solutions for ABCCB and items D, E, F, G, H with values of $a(D) = 1$, $a(E) = 2$, $a(F) = 3$, $a(G) = 4$, $a(H) = 5$, $b(D) = 3$, $b(E) = 1$, $b(F) = 5$, $b(G) = 2$, $b(H) = 4$, and the third player's values $c(D) = 4$, $c(E) = 3$, $c(F) = 1$, $c(G) = 2$, $c(H) = 5$.

CHAPTER 10

Example: Shubik Auction I

Prerequisites: Chapters 1 and 8.

You are probably familiar with English auctions, where players bid for an item. The one with the highest bid gets the item, and pays his or her bid for it. There are many versions dealing with the details, for example whether bidding increments are required, or whether the players must bid in a special order (usually they do not have to).

An English auction is easy to analyze. A player bids as long as the bid is below the worth of the item to him or her, but does not go above that.

In an attempt to make more money for the item, the auctioneer may impose the following rules:

SHUBIK AUCTION(A, B, n) Two players, Ann and Beth, bid sequentially for an item. The bids must increase in increments of 10 units. Ann starts by bidding 10 units or passing (in which case Beth can get the item for 10 units, or can instead pass as well). If Ann bids 10 units, Beth can bid 20 units or pass, and so on. After one player passes, the other player gets the item for her highest bid, but in contrast to ordinary auctions, the other player still has to pay her highest bid to the auctioneer, but gets nothing in return. There is a maximum number n of rounds. The item has a worth of A units for Ann and of B units for Beth.

Why do we need a maximum number of rounds? Without it, we would not have a finite game—the bidding could go on and on. Finiteness was also a requirement in Zermelo’s Theorem—backward induction cannot be applied to infinite games. Limiting the number of rounds is reasonable—at some point any auction should conclude.

Student Activity Play SHUBIK AUCTION(100, 100, 50) by auctioning an item that is worth a dollar (100 cents) to two students, with increments of 10 cents, with a maximum of 50 rounds. The winner wins the item for a multiple of 10 cents. At the end of the game the loser has to pay her highest bid as well.

On first sight, this looks advantageous for the auctioneer. Since she or he gets the sum of the highest and second highest bid, shouldn’t she or he make more money than with an ordinary English auction? This first impression may also have been confirmed by the class experiment. Usually, when the required end after the 50 rounds isn’t emphasized, people bid and bid until both are above the worth of the item for them. But this impression will not be confirmed by the analysis. It turns out that players should play differently, and the optimal play is not obvious.

In Figure 10.1, the game tree for SHUBIK AUCTION(25, 35, 9) is shown (the item is worth \$25 to Ann and \$35 to Beth, and the maximum number of rounds is 9).




Figure 10.1. Game tree for SHUBIK AUCTION(25, 35, 9)

Using the game tree and backward induction, any version of the game can be analyzed. But we can analyze all versions of the game at once.

Last move: As usual with backward induction, we start with the position having only end positions as successors. This occurs when the player to choose in the last round decides whether to bid or to pass. This player could be either Ann or Beth, depending on whether the number of rounds is even or odd. Let's call her the end player. Since the difference in what she pays is 20—the end player pays $10 \cdot n$ if she bids and $10 \cdot (n - 2)$ if she passes—and since she gets the unit in one case but not the other, the end player would bid in the last round provided the unit is worth at least 20 units to her, which is the case here.

Second to last move: In the round before the last round, the player to decide there (the non-end player) knows that the end player would bid in the last round and get the item. If the non-end player would raise her bid, she eventually would have to pay 20 units more, so she would not do it.

Third from last move: In the round before that, the end player has to decide. Knowing that the non-end player would pass in the next round, the end player would raise her bid.

Proceeding in this way, the conclusion is that the non-end player will always pass, whereas the end player is always prepared to raise her bid, if the item is worth at least 20 units to the end-player. Depending on the parity of the number n of rounds, we have two cases:

- If n is even, then Beth is the end player. Then Ann always passes, even in the first round, hence Beth gets the unit for 10, with a net gain of $B - 10$, and Ann, who didn't bid at all, doesn't have to pay anything and has a payoff of 0.
- If n is odd, then Ann is the end player. Ann bids 10 units in the first round, and then Beth passes. Ann's payoff is $A - 10$, and Beth's payoff is 0.

The case where the item is worth less than 20 units for a player is also easy to analyze. The player would never raise her bid. Knowing that the other player would bid in the next round provided that other player values the item at more than 20 units, the low-value player would not bid in the first round. But if both players value the item less than 20 units, the player starting would bid and get the item provided it is worth at least 10 units to her.

Therefore this auction is not good for the auctioneer provided the players play optimally. The auctioneer gets nothing for even n , and 10 units for odd n , independent of how much the item is worth (we assume it is worth at least 20 units to the end-player). However, in real life, the outcome might be different. Real people tend to be willing to raise longer. Perhaps the reason for this is that in real time players may not analyze the game completely. They do not see clearly the asymmetry of the game and the advantage for the end player.

The Wikipedia article [Wn.d.] on the game describes a variant without a limitation on the number of moves, and calls it a paradox. If we look at this game, nothing really happens. If two players start bidding, then they will continue bidding, and bids will get larger and larger. But no harm is actually done, since the players will continue to bid forever and will never have to pay.

Exercises

1. Draw the game tree for SHUBIK AUCTION(35, 35, 7), and analyze the game using backward induction.
2. Draw the game tree for SHUBIK AUCTION(5, 15, 7), and analyze the game using backward induction.
3. Draw the game tree for SHUBIK AUCTION(35, 35, 8), and analyze the game using backward induction.
4. What happens if the player not having bid last (and therefore not having gotten the item) has to pay only half of his or her last bid? Draw the game tree if the item is worth 35 for both players, and there are at most 8 rounds, and analyze the game using backward induction.
5. Discuss informally, without using a game tree, the variant where there is a number of players, bidding can be done at any time by everybody, with bids increasing by 10 units, the last two bidders have to pay their bids, and the last bidder gets the item. Assume there is a limit of 50 rounds, and that the item has a worth of 35 for each bidder.

Project 17

SHUBIK AUCTION(A, B, n) with values decreasing over time Suppose the value of the item for each player decreases by 10% each round. The item is aging. If the initial values for Ann and Beth are 25 and 35, they are 22.5 and 31.5 after the first round, 20.25 and 28.35 after the second round, and so on. Analyze the game for values 25, 35, and game length 9, and for game length 10. What about larger game lengths? What can be said for other start values than 25 and 35?

CHAPTER 11

Example: Sequential Doctor and Restaurant Location

Prerequisites: Chapters 1, 2, 8, 4, and 5.

In real life, where rules are not always carved in stone, a player in a simultaneous 2-player game may be able to move a little earlier than the other, or delay slightly until he or she sees what the other player played. If moving earlier is not possible, sometimes you can announce the move you will play, to make a commitment. If the other player believes you, it is as if you have moved first. The simultaneous game is transformed into a sequential game.

There are two roles in the sequentialization of a two-player game, moving first or moving second, and which is preferred is often decided by the temperament of the players. Some like to move first and dictate the action, and others prefer to wait and see what the other has played, concealing their intentions. However, whether it is better to move first or last should not be left to temperament but to an analysis of the game.

In this chapter we discuss the sequential versions of the location games whose simultaneous versions have been discussed in Chapters 4 and 5. In the sequential version, Ann chooses a location first, and after that Beth decides where to put her location.

Will the results differ from those in the simultaneous version? Who will gain from having to move sequentially?

11.1 General Observations for Symmetric Games

Both simultaneous location games are **symmetric**—both players have the same options, and the games are fair insofar as if both players switch their options, then the payoffs will also switch. However, the pure Nash equilibria do not necessarily have the same payoffs for both players. Still, if there is a Nash equilibrium with a payoff of a for Ann and b for Beth, then there is a Nash equilibrium with reversed payoffs of b for Ann and a for Beth. Because of this, whether the sequential version of a symmetric game gives the first mover an advantage is easy to determine—we compare both payoffs—first moving Ann’s and second moving Beth’s—in the backward induction outcome of the sequential version. If Ann’s payoff is higher, we say the game has a *first mover advantage*. If the payoffs are equal, we call the sequential version of the game *fair*. If Beth’s payoff is higher we say the game has a *second mover advantage*.

Backward induction analysis of the sequential version carries over to bimatrix notation easily: If Ann moves first and Beth second, then Ann determines the row with her move. Each row corresponds to a situation Beth may face, and in each row, Beth would choose the cell maximizing her payoff. This is her best response to Ann’s move. We color all these cells. The colored cells include all pure Nash equilibria, since the condition for coloring is only one of two required for pure Nash equilibria (both moves being best response to the other). Also, every row contains at least one colored cell. We assume that Beth is indifferent

to Ann, meaning that if she faces a position—a row—with several outcomes with the same maximum payoff, she will select any of these. So in each row, any of the colored cells could occur as the outcome. If on the other hand Beth is slightly hostile or slightly friendly to Ann, we would color only the corresponding outcomes in each cell. We finish the backward induction as follows: Since Ann knows how Beth would react in each position, Ann will choose the row giving her the largest payoff. If every row contains just one colored cell, she would select the row where Ann’s payoff in the colored cell is largest. If some cells contain multiple colors, Ann would look at the averages of her payoffs in the colored cells in each row, and compare them.

Example 1 In the classical formulation of the game CHICKEN, two cars drive towards each other. We assume that it is a simultaneous game so at some time, both have to decide simultaneously whether to proceed or to swerve. The one swerving loses face. If none swerves, both lose their lives.

	swerve	not
swerve	-1, -1	-2, 2
not	2, -2	-100, -100

This symmetric game has two pure Nash equilibria, where one swerves and the other doesn’t. In the analysis of the sequential version, we color some cells as described above. We see a first mover advantage—Ann moving first would not swerve, therefore Beth must, and the payoffs achieved are 2 and -2. If your opponent publicly removes his or her steering wheel before you start, leaving him or her no possibility to avoid a collision, then you will certainly chicken and look bad! Then the other player changed the game into a sequential game with him or her moving first.

Example 2 Our second example is ROCK-SCISSORS-PAPER. There is not much mathematical analysis necessary to see that the sequential version of this game has a second mover advantage.

Why do these two games behave so differently? Both are symmetric and both are zero-sum. In both, there is only one best response to a move—in every row, only one cell is colored in the backward induction analysis. Why does CHICKEN have a first mover advantage and ROCK-SCISSORS-PAPER a second mover advantage?

11.2 Doctor Location

Student Activity Play the sequential version of doctor location in applet [LocationDr30](#).

The game seems to have a second mover advantage. No matter where Ann moves, Beth can always get a payoff of at least 5. Since the game is constant-sum, Ann will therefore achieve at most 4.

Here is the payoff bimatrix, with Beth’s best responses to Ann’s moves underlined.

	1	2	3	4	5	6	7	8	9
1	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	4, <u>5</u>	4.5, 4.5	6, 3	5.5, 3.5	5, 4
2	4.5, 4.5	4.5, 4.5	4.5, 4.5	4, <u>5</u>	4.5, 4.5	5, 4	6, 3	5, 4	5.5, 3.5
3	4, <u>5</u>	4.5, 4.5	4.5, 4.5	4.5, 4.5	5, 4	4.5, 4.5	5, 4	6, 3	5.5, 3.5
4	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5	4.5, 4.5	4, <u>5</u>	5.5, 3.5	6, 3	5, 4
5	5, 4	4.5, 4.5	4, <u>5</u>	4.5, 4.5	4.5, 4.5	4.5, 4.5	5.5, 3.5	5, 4	6, 3
6	4.5, 4.5	4, <u>5</u>	4.5, 4.5	5, 4	4.5, 4.5	4.5, 4.5	5, 4	5.5, 3.5	6, 3
7	3, <u>6</u>	3, <u>6</u>	4, 5	3.5, 5.5	3.5, 5.5	4, 5	4.5, 4.5	4.5, 4.5	4.5, 4.5
8	3.5, 5.5	4, 5	3, <u>6</u>	3, <u>6</u>	4, 5	3.5, 5.5	4.5, 4.5	4.5, 4.5	4.5, 4.5
9	4, 5	3.5, 5.5	3.5, 5.5	4, 5	3, <u>6</u>	3, <u>6</u>	4.5, 4.5	4.5, 4.5	4.5, 4.5

Ann will choose the row for which her average highlighted payoff is highest. This is 4 in the first six rows, and 3 in the rest, so she will choose any of the first six rows. Therefore Ann’s payoff in the sequential version is 4 and Beth’s 5; there is a second mover advantage.

The simultaneous version on the graph does not have pure Nash equilibria. Since producing payoff bimatrices is tedious, from now on we will use the Excel sheet `LocationGame.xlsx`. It produces the payoff bimatrix for the doctor location game and for the restaurant location game belonging to the graph; it also calculates pure Nash equilibria for both versions and the backward induction solutions of the sequential versions of the games.

Student Activity Use the Excel sheet to check whether there is a first or second mover advantage in the doctor location game played on the graphs in Figure 11.1. Check whether there are pure Nash equilibria. Do the same for two 8-vertex graphs.




Figure 11.1. Two graphs

Why is there never a first mover advantage in such doctor location games? Since they are constant-sum, and Beth can always get half of the possible total payoff by placing her doctor in the same town as Ann. Thus Beth can always achieve as much as Ann, and sometimes more.

Can we tell from the structure of the graph which games are fair and which have a second mover advantage? Is there a feature of the graph that implies fairness or second mover advantage? Maybe, but I am not aware of any such characterization. If we stop staring at the graph, and instead look at the Nash equilibria, we can answer the questions, not only for doctor location graphs but also for sequential versions of arbitrary constant-sum symmetric simultaneous games. Proceed to the next section.

11.3 Constant-Sum Games

In sequential versions of constant-sum games, Ann just gets her security level in the simultaneous version, since the colored cells in each row—Beth’s highest payoffs in each row, Beth’s best responses—are the moves

of minimum payoff for Ann in each row. So the coloring of cells in rows and selecting the row with the best outcome of the colored cell for Ann corresponds to the maximin procedure in simultaneous games.

We have seen in Chapter 2 that a zero-sum game has a pure Nash equilibrium if the security level is 0, i.e., if the fair payoff of 0 for both can be achieved by Ann’s maximin strategy. Similarly, constant-sum games with constant sum c have a Nash equilibrium precisely when Ann’s security level reaches $c/2$. This means, by the remarks above, that the sequential version of the game is fair.

Fact *The sequential version of a symmetric simultaneous constant-sum two-person game is fair if there is some pure Nash equilibrium, and has a second mover advantage otherwise.*

This is the reason why sequential ROCK-SCISSORS-PAPER has a second mover advantage.

11.4 Restaurant Location

The Excel sheet [LocationGame.xlsm](#) can also be used for restaurant location games.

Student Activity Use the Excel sheet to check whether there is a first or second mover advantage in the restaurant location game played on the graphs in Figures 11.2 and 11.3. Investigate two more graphs with seven to ten vertices.




Figure 11.3. Graph 4

Graph 11 in Figure 11.2 was analyzed in Chapter 5. It has four Nash equilibria: Vertex 1 versus vertex 5, vertex 5 versus vertex 6, and both reversed cases. In the sequential version, Ann will play vertex 5, and Beth will play vertex 1 or 6. This way, Ann gets a payoff of 2.5 and Beth a payoff of 2. There is a first mover advantage.

Graph 4 in Figure 11.3 has six Nash equilibria: Vertex 2 versus vertex 4, vertex 5 versus vertex 8, vertex 6 versus vertex 9, and the reverse cases. In the sequential version, Ann plays vertex 2 or 4, and Beth plays vertex 4 or 2, respectively. The game is fair—both players get a payoff of 3.

We have seen that sequential versions of doctor location games never have a first mover advantage. Based on the examples, couldn’t it be just the opposite—never a second mover advantage—for restaurant location games?

Conjecture Sequential versions of restaurant location games never have a second mover advantage.

Student Activity In the applets [LocationGrid8seq](#) and [LocationGrid16seq](#) you can create various graphs and perform backward induction analysis. You can also play against the computer. Check the

conjecture for about ten graphs. In the same applets you can also check whether versions with two restaurants for each chain never have a second mover advantage.

11.5 Nash Equilibria and First Mover Advantage for Symmetric Games

We know that all simultaneous restaurant location games have pure Nash equilibria (see Chapter 5). We suspect, based on several examples, that restaurant location games never have a second mover advantage. Could there be some connection? We showed that constant-sum symmetric games with pure Nash equilibria have no second mover advantage (since they are fair), which fits into the picture we have painted so far. So, is it true that all symmetric games with pure Nash equilibria have no second mover advantage? If this were true, the conjecture would follow as a special case. We try the same approach as before, moving to a larger class of games, from restaurant location games to symmetric games with pure Nash equilibria.

It turns out that the conjecture is true under one additional condition: we assume that there is only one best response to each one of Ann’s moves, that is, in each row only one cell is colored. We call this property **unique response** for Beth. Then all pure Nash equilibria are among the colored cells, and Ann can choose any of them. So assume Ann chooses the Nash equilibrium maximizing her payoff among all Nash equilibria, and this move gives a to Ann and b to Beth. Then a is as least as large as b , since there is also a Nash equilibrium giving b to Ann and a to Beth—the one where Ann and Beth exchange moves. But Ann doesn’t have to choose a Nash equilibrium—it may be that a row with no Nash equilibrium in it promises more to Ann than a . Then Ann would choose this row, but still get more than Beth.

Fact *If a symmetric simultaneous 2-player game has pure Nash equilibria and Beth’s response to every move of Ann is unique, then the sequential version of the game is fair or has a first mover advantage.*

Often the first condition fails. People are often indifferent between options. For restaurant location games we rarely have this property.

Exercises

1. Look at the simultaneous 2-player game described by the payoff matrix

	M1	M2
M1	1, 1	4, 2
M2	3, 5	3, 3

Is it symmetric? Does it have a Nash equilibrium? Who is better off in the sequential version, with the row player (Ann) moving first?

2. Look at the simultaneous 2-player game described by the payoff matrix

	M1	M2	M3
M1	1, 1	3, 4	1, 3
M2	4, 3	3, 3	1, 5
M3	3, 1	5, 1	2, 2

Is it symmetric? Is it constant-sum? Does it have a Nash equilibrium? Is there a first mover or a second mover advantage in the sequential version, with the row player (Ann) moving first?

3. Look at the simultaneous 2-player game described by the payoff matrix

	M1	M2	M3
M1	3, 3	2, 4	5, 1
M2	4, 2	3, 3	2, 4
M3	1, 5	4, 2	3, 3

Is it symmetric? Is it constant-sum? Does it have a Nash equilibrium? Is there a first mover or a second mover advantage in the sequential version, with the row player (Ann) moving first?

4. Look at the simultaneous 2-player game described by the payoff matrix

	M1	M2	M3
M1	0, 0	2, -2	1, -1
M2	-2, 2	0, 0	2, -2
M3	-1, 1	-2, 2	0, 0

Is it symmetric? Is it constant-sum? Does it have a Nash equilibrium? Is there a first mover or a second mover advantage in the sequential version, with the row player (Ann) moving first?

- Analyze the sequential version of restaurant location in the left graph in Figure 11.1, under the assumption that in case of ties, options are chosen with equal probability.
- Analyze the sequential version of restaurant location in the right graph in Figure 11.1, under the assumption that in case of ties, options are chosen with equal probability.
- Use the Excel sheet [SIMSEQsym.xlsx](#) to test whether symmetric games with Nash equilibria have no second mover advantage. Generate ten such games and check.

Project 18

Try to reject the conjecture stated in Section 11.4. Create ten arbitrary restaurant location games on graphs with seven to ten vertices, and check whether they have no second mover advantage using the Excel sheet [LocationGame.xlsm](#). Use the applets [LocationGrid8seq](#) and [LocationGrid16seq](#) for more special graphs. Discuss whether this makes the conjecture true, or maybe more likely to be true.

Project 19

Hostile versus Friendly Play We saw in Subsection 8.3.4 that there are several ways to break ties in backward induction. Applets [LocationGrid8seq](#) and [LocationGrid16seq](#) use the neutral version. Use applets [LocationGrid8seq3](#) and [LocationGrid16seq3](#) to compare the outcomes in optimal play for all three versions, for various graphs generated, for the 2-restaurant and for the 4-restaurant versions. Are the payoffs in the friendly versions always higher than in the neutral versions? Who seems to profit most from friendliness? Investigate many games in these applets, and discuss your findings.

CHAPTER 12

Theory 4: Probability

Often games contain random features. In poker and other card games, the first move is made not by the players but by the dealer who deals the cards randomly. Outcomes depend on the players' moves and on random elements. Even some completely deterministic games, like ROCK-SCISSORS-PAPER, are best played using some random device. Accordingly, we need to discuss the theory of probability.

12.1 Terminology

Much mathematical reasoning is concerned with trying to predict outcomes. If I create a sphere of radius 20cm, how large will its surface area be? If I combine 20ml of a 20% acid solution and 30ml of a 50% acid solution, what is the strength of the solution I get? What will be the speed of a falling apple 0.1 seconds after it is dropped? If a roulette ball has an initial speed of 3 meter per seconds, a given initial direction, and the roulette wheel is spinning at a certain speed, where will the ball end? If interest rate increases by 5% and productivity by 3%, how will the unemployment rate change? In many situations, models from the sciences, social sciences, or economics are not strong enough to predict these outcomes. Or, the models may be accurate, but the data available is not sufficient, which is the case in the roulette example. In those cases, outcomes seem random or unpredictable.

This is where probability theory enters. There is usually a better prediction than “I don't know what the outcome will be, since it is random anyway.” Although you cannot predict the outcome, you may be able to describe the likelihood of possible outcomes. This may tell you what you may expect. Expectations are not always realities, but we will see that in the long run they will be.

Probability theory uses some terminology that we have to get used to. For instance, the situations whose possible outcomes are to be predicted are called **experiments**. Ideally, like scientific experiments, these experiments are not singular and isolated, but can be repeated since their settings and conditions are very clearly defined. Examples are rolling a die, throwing a coin, or selecting a card from a shuffled 52-card deck. The weather tomorrow or the outcome of elections could also be called experiments. We concentrate on one feature of an experiment, for example what number shows on top of a die after it comes to rest, or what card is chosen, and neglect other features (as, for instance, how long the die rolls). When describing an experiment, you should always describe the procedure and specify what feature of the outcome you are focusing on.

Often only finitely many outcomes are possible in an experiment. They are called the **simple events**. **Events** are groups of simple events. For instance, each of the 52 cards is a simple event in the experiment of drawing a card from a shuffled deck of cards. Examples of events are drawing an ace (combining four simple events) or drawing a diamond (combining 13 simple events), but an event does not have to be expressible as elegantly as in these two examples. Sometimes events can be described only by listing the simple events they consist of. For instance, drawing either a club king, or a heart 8, or a diamond 3 is also an event.

Student Activity Take as an example the experiment in DMA100, a soccer simulation. The game is discussed in Chapter 14, but for now you don't have to know the details. An experiment is performed by clicking twice on the "Step-by-Step" button. The ball ends on the left (goal for red), on the right (goal for blue), or in the middle (no goal). These are the three outcomes.

Perform the experiment 100 times. You can automate this by clicking the "Automatic Play" button, and you can speed it up by clicking the "faster" button repeatedly. You will notice that the outcome "goal for red" will occur more frequently than the outcome "goal for blue". Does this mean that a goal for red is more likely than a goal for blue?

The task of probability theory is to assign to each event A a number $p(A)$ between 0 and 1, called "the **probability** of A ". It measures the likelihood that the outcome of the experiment belongs to A . If $p(A) = 0$, it is almost always the case that A cannot occur. If $p(A) = 1$, it is almost always the case that A has to occur. Everything between these two extremes is possible.

There are different ways to obtain probabilities. One way is to perform the experiment often and measure how often the event occurs. If you perform the experiment n times, and if k times the event A occurs, then the **empirical probability** $p(A)$ of that event is defined to be the **relative frequency** k/n of the event. A disadvantage is that we have to perform the experiment repeatedly, and that the probabilities we get vary.

However, with "a priori" models we avoid having to perform the experiment. We obtain values of **theoretical probabilities** using mathematical reasoning. The values should predict the empirical (relative frequency probability) in the following sense:

Theorem (Law of Large Numbers) *As an experiment is repeated, the empirical probability (relative frequency) of an event will converge to its theoretical probability.*

If not, the a priori model is not a good reflection of reality.

Student Activity In the DMA soccer example, if you don't move the soccer players, the theoretical probabilities for a goal for red and blue are $1/3$ and $1/5$ (as will be shown in Chapter 14 using tools from Section 12.4). The empirical probabilities you obtained in the student activity might have been close to these numbers. If you calculate the averages of averages obtained by all students, they are very likely (according to the Law of Large Numbers) to be closer to $1/3$ or $1/5$ than most of the individual averages. Check it out!

Let's assume that the theoretical probability for the event head in the experiment of flipping a coin is $1/2$. Assume that frequencies of head after 100, 1000, and 10000 throws are 49, 550, and 5200. Then the corresponding relative frequencies are $49/100 = 0.49$, $550/1000 = 0.55$, and $5200/10000 = 0.52$. The Law of Large Numbers does not say that the relative frequencies cannot temporarily move away from 0.5, as from the close 0.49 to the disappointing 0.55. What it says is that almost surely eventually the relative frequency will come closer to 0.5. Another way to say this is that a deviation of, say, 0.03 (a value below 0.47 or above 0.53) is possible, but its probability decreases as the number of trials increases. Actually the probability approaches 0. However, the absolute deviation can always exceed a fixed value. In the example, the absolute deviation between the number of heads and the predicted number of heads is -1 , 50, and 200, and these numbers may increase as the number of trials increases.

Let me mention a common misconception about the Law of Large Numbers. You throw a coin and get a sequence of head, tail, tail, head, tail, head, tail, tail, tail, tail, tail. How likely is "head" next? Many people would say it is larger than 50%, since it is due. Others may question the fairness of the coin and rather predict another tail. However, if we know that we have a fair coin, then the theoretical probability of 50% is the right one and the probability for head is still exactly 50% for all future rounds. Probability theory and the Law of

Large Numbers are not fair in the way of fighting against an existing deviation. The deviation may persist, or even increase, but then the relative deviation will still go to 0.

Historical Remark Although Gerolamo Cardano did some work on probabilities in the 16th century, it was not much noticed, and probability theory got its real start in 1654 in an exchange of letters between Blaise Pascal and Pierre de Fermat, as described in a separate chapter. Christian Huygens published the first book on this new theory of probability in 1657. The Law of Large Numbers was formulated and proved by Jakob Bernoulli in 1713. One hundred years later, Pierre de Laplace introduced new tools, and applied probability theory to areas outside of games of chance. Later the area of statistics emerged from probability theory.

12.2 Computing Probabilities

12.2.1 Equally Likely Simple Events

The probabilities of all simple events always sum to 1, since the event containing all simple events occurs. Therefore it is simple to find probabilities when it is known that all outcomes are equally likely (as for a fair die). If there are n possible outcomes, n simple events, then the probability of each is $1/n$.

The probability $p(A)$ of an event A is the sum of the probabilities of the simple events contained in it. If all simple events have the same probability then $p(A)$ is the number of simple events in A divided by the total number of simple events.

Example 1 How likely is it to throw two fives with two dice of different color?

We roll two dice, a red one and a blue one. The simple events are the 36 pairs of numbers $(1, 1), (1, 2), \dots, (6, 5), (6, 6)$, where the first entry indicates what the red die shows and the second entry what the blue die shows. A red 3 and a blue 5, abbreviated as $(3, 5)$, and a red 5 and a blue 3, denoted by $(5, 3)$, are different outcomes. All the simple events are equally likely, so each one has a probability of $1/36$. Thus $p((5, 5)) = 1/36$.

12.2.2 Simple Events not Equally Likely

More difficult is the situation when the outcomes are not equally likely, as for a biased die. Sometimes one has to use the empirical probability $p(A)$, which is defined as the relative frequency. Sometimes it is possible to redefine these simple events so they are equally likely. For example

Example 2 A container contains 5 red marbles, 4 blue marbles, and 3 green marbles. Closing your eyes, you remove one randomly. What is the probability of picking a red marble?

If we define red, blue, and green as simple events, then they are not equally likely. It is not difficult to assign probabilities to them, but we can formulate the problem differently. Instead of having three simple events, we could label the different marbles and have 12 different equally likely simple events: $\text{red}_1, \text{red}_2, \dots, \text{green}_2, \text{green}_3$. The event red consists of the five simple events $\text{red}_1, \dots, \text{red}_5$, therefore $p(\text{red}) = 5/12$.

Example 3 How likely is it to throw two fives with two dice?

We could look at pairs of numbers as simple events. If the dice are indistinguishable, the outcome showing 1 and 2 is the same as the one showing 2 and 1. Then there are 21 possible outcomes, but they are not equally likely. It is simpler to color one die. Then we have 36 equally likely outcomes with

probability $1/36$ each. Therefore $p(5 \text{ and } 5) = 1/36$, but $p(1 \text{ and } 2) = 2/36$, since two simple events in the colored version would correspond to the simple event of one die showing a 1 and the other showing a 2.

Example 4 How likely is it to get a sum of four with two dice?

It would be possible to formulate the simple events as the sums of the numbers shown by the two dice. This model would have the disadvantage that the probabilities of the simple events would not be identical (think about how likely a sum of 2 is compared with a sum of 7). A better model uses the 36 pairs of numbers as simple events. Then the event A of getting a sum of four consists of three simple events, namely $(1, 3)$, $(2, 2)$, and $(3, 1)$. Therefore $p(A) = 3/36$.

12.3 Expected Value

Example 5 You have the opportunity to throw a coin. If it shows a head, you get 100 dollars. If it shows tail, you get 50 dollars. If it stands on its edge, you have to pay 20 cents. Would you decline the offer to play since you may lose 20 cents? Would your willingness to play change if you would have to pay 10,000,000 dollars if the coin stands on its edge?

Often numerical values are attached to outcomes of an experiment. We then have a **random variable**, abbreviated by a letter, say X . In our games the payoff to a player a random variable. If the simple events (possible outcomes) of an experiment are $A_1, A_2, \dots, A_n$ and the payoffs (the amounts attached to them) are $X_1, X_2, \dots, X_n$, then how high a payoff would you expect? Probably the average of all the values, provided the outcomes are equally likely, i.e., if $p(A_1) = p(A_2) = \dots = p(A_n) = 1/n$. The average equals the sum of the products $p(A_i) \cdot X_i$ of probabilities and amounts attached. If the outcomes are not equally likely, the expectation would be tilted towards the payoffs attached to the more likely outcomes. Instead of an average, we take a weighted average, where the probabilities of the outcomes are the weights. The expected value of the random variable X is defined by

$$E(X) = p(A_1) \cdot X_1 + p(A_2) \cdot X_2 + \dots + p(A_n) \cdot X_n.$$

Just as the relative frequency of an outcome approaches its probability if an experiment is repeated many times, the average value of a random variable approaches its expected value when an experiment is repeated many times.

Example 6 Three amounts of money, \$10, \$100, and \$1000 are in three closed identical envelopes. One of them is randomly given to you. How much money do you expect?

Your expectation is larger than the middle value, \$100. Since the envelopes are equally likely, with probability $1/3$, the expected value is $(1/3) \cdot 10 + (1/3) \cdot 100 + (1/3) \cdot 1000 = (1/3) \cdot 1110 = 370$. We expect a value that doesn't occur among the values inside the envelopes. But if we play the game 1000 times, we would expect to average about \$370 in each play.

The expected value of a random value is not necessarily the value that occurs most frequently or the most likely value.

Example 7 You draw a card at random from a deck of 52 cards. You win \$10 if you choose an ace, \$3 for a king, and nothing otherwise. What are your expected winnings?

We have 52 simple events. To the four aces we attach \$10, to the four kings \$3, and to the remaining 44 cards \$0. Each of the 52 simple events has probability $1/52$. Accordingly your expected winnings are $44 \cdot (1/52) \cdot 0 + 4 \cdot (1/52) \cdot 3 + 4 \cdot (1/52) \cdot 10 = 52/52 = 1$ dollar.

Example 8 In a European roulette game, all numbers between 0 and 36 are equally likely. If you bet on even, you obtain \$1 if the result is one of the 18 even numbers larger than 0. Otherwise, if the result is odd or 0, you have to pay \$1. What are your expected winnings?

All 37 simple events are equally likely with probability $1/37$. 18 of them have the value 1 attached, 19 have the value -1 attached. Then the expected value is $18/37 - 19/37 = -1/37$.

Example 9 We roll two dice and consider the sum of the two dice as a random variable. What is its expected value?

The most convenient way to model this is by assuming the dice have different colors, and we therefore have 36 different possible outcomes. One simple event has 2 as the sum of the dice, two have sum 3, three have sum 4, and so on until sum 7, which occurs six times. From then on, the frequency for the sum decreases. We have 5 cases of a sum of 8, 4 cases of a sum of 9, and so on, until the one simple event of a sum of 12. The expected sum of the two dice is

$$\begin{aligned} \frac{1}{36} \cdot 2 + \frac{2}{36} \cdot 3 + \frac{3}{36} \cdot 4 + \frac{4}{36} \cdot 5 + \frac{5}{36} \cdot 6 + \frac{6}{36} \cdot 7 + \\ + \frac{5}{36} \cdot 8 + \frac{4}{36} \cdot 9 + \frac{3}{36} \cdot 10 + \frac{2}{36} \cdot 11 + \frac{1}{36} \cdot 12 = 7. \end{aligned}$$

Expected values can be computed only for numerical experiments. If you draw a marble from a container containing six blue marbles and six yellow ones, the expected outcome is not a green marble. There is no such thing as an expected outcome. In game theory we have expected payoffs. Payoffs are random variables.

Modeling Note Measurement Scales Assume there are three possible outcomes: bad, good, and very good, and let's assume we translate these outcomes into payoffs of 0, 1, and 2. Then would you prefer a 60% chance of a very good outcome and a 40% chance of a bad outcome to getting a good outcome all the time? The arithmetic is easy: $60\% \cdot 2 + 40\% \cdot 0 = 1.2 > 1$, so you would. This assumes that the assignment of 0, 1, and 2 to the outcomes is appropriate. If the difference between very good and good is small, but the difference between good and bad is large, risking a bad outcome in 40% of the cases would be reckless.

There are four types of measurement scales: nominal, ordinal, interval, and cardinal. A **nominal scale** is the weakest. Some experiments' outcomes have no natural order, like red, green, and blue. Outcomes are **ordinal** if there is a natural ordering. An example is bad, good, and very good. Ordinal outcomes can be encoded with numbers. Outcomes on an ordinal scale are on an **interval scale** if they can be encoded with numbers such that the differences between the numbers are meaningful and comparable. In the example, the outcomes are interval if we can tell whether good is closer to bad or to very good, and to what extent. If good is closer to very good than to bad, then encoding bad, good, and very good by 0, 1, 2 may not be appropriate, but encoding them by 0, 2, 3 may be appropriate. An interval measurement scale is a **cardinal scale** or **ratio scale** if there is an interval encoding of the outcomes with a natural (meaningful) zero.

For game theory, we attach payoffs for each player to all outcomes. To do this, the outcomes of a game should be at least ordinal. Players must be able to order the outcomes from least favorable to most favorable. Then we could assign payoffs of 1, 2, ... to them, from least favorable to most favorable. Trying to find the best outcome for a player means finding an outcome with maximum payoff.

If the outcomes are not on an interval scale, then a player would not necessarily be able to tell whether he or she prefers a 50%-50% chance of outcome A or B to a sure outcome of C. Assume that Ann’s payoffs for outcomes A, B, and C are 2, 5, and 3. Only with an interval scale would the expected value of $0.5 \cdot 2 + 0.5 \cdot 5 = 3.5$ imply that indeed the player prefers the mix to outcome C. All considerations comparing expected payoffs require interval scales.

Outcomes of games are often on a ratio scale. Examples are games where the payoff is money won or lost, where positive means winning and negative means losing. Then a payoff of 0 divides winning from losing.

12.4 Multistep Experiments

12.4.1 Probability Trees

Some experiments consist of two or more sub experiments that are performed sequentially, called multistep experiments. For example we first have an experiment with outcomes A and B. If the outcome is A, a second experiment is performed with outcomes C and D. If the outcome of the first experiment is B, a second experiment with outcomes E and F is performed, see Figure 12.1, a **probability tree**. Starting with the leftmost vertex, one of the two vertices in the middle is reached after the first experiment, and then one of the four vertices on the right. The probabilities for the outcomes are written below the lines leading to the outcome. There is a similarity to extensive forms or game trees in sequential games. What was called position there is here called situation—the combination of all that is known so far.




Figure 12.1. A probability tree

How likely is outcome A in the first round and outcome C in the second round? The probability for the combined event is the product of the probabilities on the arcs, $p(A) \cdot p(C)$, assuming that the experiments are independent, that is, that the outcome of the first experiment does not influence the probabilities in the second experiment.

12.4.2 Conditional Probabilities

Often in multistep experiments, the outcomes of the second-round experiments are the same, but their probabilities differ depending on the outcome of the first-round experiment. An example is shown in Figure 12.2. Since more than one arc in the tree has C as outcome, but the probabilities may be different, we need a more involved notation. Let $p(C|A)$ denote a **conditional probability**, the probability for outcome C provided we are in the situation generated by outcome A in the first round. The probabilities $p(C|A)$ and $p(C|B)$ may not be identical. For the combined experiment, four outcomes are possible: A and C (A&C), A and D (A&D), B and C (B&C), and B and D (B&D). How likely is each? The total probability for the combined outcomes on the right are the products of the probabilities of the path leading to it.




Figure 12.2. A probability tree with duplicate outcomes

That is

$$p(A \& C) = p(A) \cdot p(C|A),$$

$$p(A \& D) = p(A) \cdot p(D|A),$$

$$p(B \& C) = p(B) \cdot p(C|B),$$

$$p(B \& D) = p(B) \cdot p(D|B).$$

Example 10 AIDS tests are not perfect. Assume your hospital uses one that gives a positive result in 97.7% of the cases if the person has the AIDS virus, and a negative result in 92.6% of the cases if the person is not infected. Assume that 0.5% of the population in your city carries the AIDS virus.

If someone from your city tells you that he or she has tested positive, what is the probability that he or she carries the virus? [RC2004]




Figure 12.3. Aids test

We model this as a two-step experiment, see Figure 12.3. In the first round, the outcomes are AIDS infected (A) with probability 0.5% and not AIDS infected (H for healthy) with probability 99.5%. In the second step, the test is positive (P) or negative (N) with probabilities depending on the outcome of the first step. The first and third outcomes of the experiment combined are the outcomes of a positive test. The first outcome is that the person has AIDS and tested positive. Its probability is $p(A) \cdot p(P|A) = 0.005 \cdot 0.977 = 0.004885$. The third outcome is that the person is not AIDS infected but tested positive. Its probability is $p(H) \cdot p(P|H) = 0.995 \cdot 0.074 = 0.07363$. Then the probability that the person who tested positive is infected is $\frac{0.004885}{0.004885 + 0.07363} = 6.34\%$. We see that testing positive is not a reason for panic.

Example 11 A deck of 52 cards is available. The first experiment consists of drawing a card with outcomes red or black. Let R_1 be the outcome of drawing a red card, and B_1 be the outcome of drawing a black card. Then another card is drawn from the remaining 51 cards. Let R_2 be the event that it is red, and B_2 that is black.

We have $p(R_1) = p(B_1) = 26/52$. Then $p(R_2|R_1)$, the probability for a second red card if one red card has already been drawn, is $p(R_2|R_1) = 25/51$. Also $p(B_2|R_1) = 26/51$, $p(R_2|B_1) = 26/51$, and $p(B_2|B_1) = 25/51$. Therefore the combined event of two red cards $R_1 \& R_2$ has the probability $p(R_1 \& R_2) = p(R_1) \cdot p(R_2|R_1) = 26/52 \cdot 25/51$. Similarly for the combined event $B_1 \& B_2$ of two black cards: $p(B_1 \& B_2) = p(B_1) \cdot p(B_2|B_1) = 26/52 \cdot 25/51$. The event $R_1 \& B_2$ of drawing a red card in the first round and a black card in the second round has probability $p(R_1 \& B_2) = p(R_1) \cdot p(B_2|R_1) = 26/52 \cdot 26/51$. We get the same probability for the event $B_1 \& R_2$ for a black card in the first round and a red card in the second round. Therefore the probability of getting cards of different colors is $p(R_1 \& B_2) + p(B_1 \& R_2) = 26/52 \cdot 26/51 + 26/52 \cdot 26/51 = 26/51$, slightly greater than $1/2$.

12.4.3 Probability Digraphs

Let's draw three cards from a standard deck, without replacement. R_1, R_2 , and R_3 mean a red card in the first, second, and third drawing, and B_1, B_2, B_3 mean a black card in the first, second, third drawing. The probability tree is in Figure 12.4.




Figure 12.4. A probability tree

Suppose we care only about how many black and red cards are drawn. Then we don't distinguish between $R_1 \& R_2 \& B_3$, $R_1 \& B_2 \& R_3$, and $B_1 \& R_2 \& R_3$ —in each case we have two red cards and one black one. After the second drawing we don't have to distinguish between $R_1 \& B_2$ and $B_1 \& R_2$. In both cases we have one red and one black card in the hand, and 25 red cards and 25 black cards remaining in the stack. So the prospects for the third drawing are identical. So that subtrees of the probability tree starting at $R_1 \& B_2$ and $B_1 \& R_2$ are identical. We will identify these states and the subtrees starting at the vertices. What we get is called a **probability digraph**. For our example it is shown in Figure 12.5.




Figure 12.5. The corresponding probability digraph

In the digraph, the probability for a combined event like 2 red cards, 1 black card would be obtained by looking at paths starting at the leftmost vertex, moving to the right, and arriving at the vertex in question. For each of the paths, multiply the probabilities along all edges, and add all these products. In the example, there are three paths, (R_1, R_2, B_3) , (R_1, B_2, R_3) , and (B_1, R_2, R_3) , leading to the outcome of two red and one black card. Therefore its probability is

$$\begin{aligned}
 & p(R_1) \cdot p(R_2|R_1) \cdot p(B_3|R_1 \& R_2) + p(R_1) \cdot p(B_2|R_1) \cdot p(R_3|R_1 \& B_2) \\
 & + p(B_1) \cdot p(R_2|B_1) \cdot p(R_3|B_1 \& R_2) \\
 & = \frac{1}{2} \cdot \frac{25}{51} \cdot \frac{26}{50} + \frac{1}{2} \cdot \frac{26}{51} \cdot \frac{25}{50} + \frac{1}{2} \cdot \frac{26}{51} \cdot \frac{25}{50} = \frac{1950}{5100} = \frac{13}{34}.
 \end{aligned}$$

Chapter 13 gives another application of this method, and contains remarks on the history of probability theory.

12.5 Randomness in Simultaneous Games

Let's return to games. Often, the payoffs in a simultaneous game depend on the choices of the players, and also have a random component. Consider the following example:

Example 12 Variant of TWO BARS: Two bars charge \$2, \$4, or \$5 for a beer. Every day 60 beers are drunk by natives, who go to the bar with lower prices. Tourists choose one of the two bars randomly. Depending on whether a tourist bus stops at the village or not (both outcomes are equally likely), 100 or 50 beers are drunk by tourists. What prices would the bars select?

Sequential games involving randomness are modeled by defining a random player, (sometimes called "Nature").

In the example, assume both bars set a price of \$2. Then both will sell 30 beers to natives. With 50% probability both will sell 50 beers to tourists, each making $2 \cdot (30 + 50) = 160$ dollars, and with 50% probability both will sell only 25 beers to tourists, each making $2 \cdot (30 + 25) = 110$ dollars. The expected value for each is $(1/2) \cdot 160 + (1/2) \cdot 110 = 135$ dollars. The same analysis could be done for the other eight entries. In this simple example we could also use the deterministic model with an expected number of $(1/2) \cdot 100 + (1/2) \cdot 50 = 75$ tourists. The modified bimatrix is

	2	4	5
2	135, 135	195, 150	195, 187.5
4	150, 195	270, 270	390, 187.5
5	187.5, 195	187.5, 390	337.5, 337.5

Example 13 A 3-spinner is a random device that shows one of three numbers with equal probability. Two players bring two 3-spinners with numbers (7, 5, 4) and (1, 8, 6) to the game, and each selects one simultaneously. The chosen spinners play against each other. The player whose 3-spinner shows the larger number wins. Which 3-spinner would you select?

If both select different spinners, there are nine outcomes possible, 7-1, 7-8, 7-6, 5-1, 5-8, 5-6, 4-1, 4-8, 4-6, and all are equally likely. In outcomes 1, 3, 4, 7, the left player wins. In the other five, the right player wins. Therefore the (1, 8, 6) spinner is better than the other one, and both will select it.

Example 14 Assume each player brings three 3-spinner to the game: (3, −1, −2), (2, 2, −4), and (−1, −1, 2). Players select one of them simultaneously, and they spin. The player whose spinner shows the higher number wins one dollar from the other. Which spinner should be selected?

This is a zero-sum game, so it suffices to display the expected payoff for the row player Ann.

	(3, −1, −2)	(2, 2, −4)	(−1, −1, 2)
(3, −1, −2)	0	1/9	−1/9
(2, 2, −4)	−1/9	0	1/9
(−1, −1, 2)	1/9	−1/9	0

It may be a little surprising to see that there are no Nash equilibria. There is no best spinner.

12.6 Counting without Counting

The probability of an event is the sum of the probabilities of the simple events contained in it. If we know the probabilities of the simple events, computing probabilities of events is simple. This is true in particular when simple events have equal probabilities. Sometimes counting simple events can be tedious. Formulas for doing this are the content of the part of mathematics called combinatorics, but we won't go into this interesting subject.

Exercises

1. Assume you flip a coin and get \$3 if the result is heads and have to pay \$2 if the result is tails. What is the expected value?
2. You have five envelopes, containing \$0, \$1000, \$2000, \$3000, and \$4000. You randomly choose one of them. What is the expected value?
3. You have five \$1-bills, three \$5-bills, and one \$10 bill in your pocket. You randomly choose one of them to give a taxi driver a tip. What is the expected value?
4. You want to buy a used car tomorrow, so you went to your bank and withdrew \$8000 in cash. In the evening there is a football game that you want to attend and you have the ticket, worth \$15. You read in the newspaper that during the last game, 5 out of 80000 spectators had their money stolen. You have lived in your apartment for 10 years and never had a burglary. What do you do? Stay at home and watch your money, or go to the game and leave the money at home, or 3) go to the game and take the money with you? Explain.
5. a) Two cards are selected randomly from a shuffled 52-card deck. Someone tells you that the first card is black. What is the probability that the second card is black?
b) Two cards are selected randomly from a shuffled 52-card deck. Someone tells you that at least one of the cards is black. What is the probability that both cards are black?

6. Analyze the following game.

SENATE RACE (compare [Kn.d.] First, an incumbent senator decides whether to run an expensive ad campaign for the next election. After that, a challenger decides whether to enter the race. The chances for the senator to win are $5/6$ with the ad campaign and $1/2$ without. The value of winning the election is 2, of losing -0.5 , and the cost of the add campaign is 1 (all values in millions of dollars).

7. a) You put \$100 on red in European roulette. If you lose, the bet is gone. If you win, you get your bet back, plus an additional \$100. How much total win or loss would you expect?
- b) You put \$10 on the numbers 1–12. If you lose, you lose the bet. If you win, you get your bet back plus an additional \$20. How much total win or loss would you expect?
- c) You put \$5 on the number 13. If you lose, you lose the bet. If you win, you get your bet back plus an additional \$175. How much total win or loss would you expect?
8. You want to insure your yacht, worth \$90,000. A total loss occurs with probability 0.005 in the next year, a 50% loss with probability 0.01, and a 25% loss with probability 0.05. What premium would you have to pay if the insurance company wants to expect a profit of \$200 per year?
9. A random number generator produces a sequence of three digits, where each of the digits 0, 1, 2, and 3 has equal probability of $1/4$. Digits are generated independently. Draw the probability tree for the three-step experiment. Find the probability that a sequence
 - a) consists of all ones,
 - b) consists of all odd digits.
10. Three digits are generated in three rounds. In the first round, a leading digit is selected. It is either 4 or 5, with equal probability. In later rounds, a digit is selected from the digits larger or equal to the previously selected one, with equal probability. That means that if the first digit selected was a 5, the next one could be one of 5, 6, 7, 8, 9, which all have probability $1/5$. If 7 is selected as second digit, then the third one must be one of 7, 8, 9, all having equal probability of $1/3$. Draw the probability tree for the experiment. Find the probability that a sequence
 - a) consists of all even digits,
 - b) consists of all odd digits.
11. There are eight balls in an urn, identical except for color. Three are blue, three are red, and two are green. Consider the three-step experiment of
 - drawing one ball, setting it aside,
 - drawing another ball, and setting it aside, and
 - drawing a third ball, and setting it aside.
 - a) Draw the probability tree or probability digraph for the experiment.
 - b) How likely is it that two of the balls drawn are blue, and one is green?
 - c) How likely is it that all three balls have different colors?
 - d) How likely is it that all three balls have the same color?
12. a) Draw a probability tree for the possible head-tail sequences that can occur when you flip a coin four times.

- b) How many sequences contain exactly two heads?
 - c) How many sequences contain exactly three heads?
 - d) Draw the probability digraph for the case where you are interested only in how many heads a sequence contains, not when they occur. That is, you would identify the situations HT and TH, you would identify HHT, HTH, and THH, and so on.
 - e) What is the probability of getting exactly two heads?
13. There are eight balls in an urn, identical except for color. Four of them are blue, two are red, and two are green. Consider the three-step experiment of
- drawing one ball, setting it aside,
 - drawing another ball, and setting it aside, and
 - drawing a third ball, and setting it aside.
- a) Draw the probability tree or probability digraph for the experiment.
 - b) How likely is it that two of the balls drawn are red and one is green?
 - c) How likely is it that all three balls have different colors?
 - d) How likely is it that all three balls have the same color?
14. A coin is flipped at most five times. When a tail occurs, the experiment is stopped. If a head occurs and if the coin hasn't been flipped more than four times, it is flipped again.
- a) Draw the probability tree of the multistep experiment.
 - b) What is the probability that exactly one tail occurred?
15. A coin is flipped at most five times. When a tail occurs the second time, the experiment is stopped. If two tails haven't occurred and if the coin hasn't been flipped more than four times, it is flipped again.
- a) Draw the probability tree of the multistep experiment.
 - b) Draw the probability digraph of the multistep experiment.
 - c) What is the probability that exactly one tail occurred?
16. You draw three cards out of a shuffled 52 cards deck. What is the probability of drawing a straight—a sequence of three consecutive cards, like 10, J, Q?
17. You draw three cards out of a shuffled 52 cards deck. What is the probability of getting three of a kind—all cards having the same rank? What is the probability of getting a pair but not three of a kind?
18. You draw three cards out of a shuffled 52 cards deck. What is the probability of getting a flush—all three cards from the same suit?

Project 20

TENNIS Steffie and Arantxa are playing a tennis set. A set consists of games and games consist of points. A game is won by the player who has first won at least four points and at least two points more than the other player. A set is won by the player who has won at least six games, and at least two more than the other player.

Assume a point is won by Steffie with probability 0.52 and by Arantxa with probability 0.48. How likely is it that Steffie wins a game? How likely is it that Steffie wins a set?

A tiebreak is the rule that if both players have won six games in a set, the player winning a tiebreak game wins the set. This special tiebreak game is won by the player who has first won at least seven points and at least two points more than the other player. How likely is it that Steffie wins a set with tiebreak rule?

Project 21

FINAL EXAM This game is played by three players: the teacher, student X and student Y. There are 20 other students, but they don't make decisions and therefore do not participate in the game. Student X makes the first move by wearing his cap to class or not. Seeing this, Student Y has the same option, cap or not. Seeing the two students with or without caps, the teacher makes the decision to keep an eye on X, on Y, or the teacher could also relax. Seeing the teacher react, students X and Y decide whether to cheat or not. Cheating, when not noticed by the teacher, brings a better grade, a payoff of 1. If the teacher notices cheating, disciplinary measures are taken, a payoff of -6 for the student. Cheating is noticed by the teacher with probability $1/5$ if the student has no cap and the teacher is not keeping an eye on the student, with probability $1/10$ if the student wears a cap but the teacher is not keeping an eye on the student, with probability $1/3$ if the student wears no cap and the teacher is keeping an eye on the student, with probability $1/6$ if the student wears a cap and the teacher is keeping an eye on the student. The teacher has also payoffs. It is the sum of a value of -1 if he or she has to keep an eye on a student, compared to 0 if he or she can relax, and a value of -2 if cheating has occurred but gone unnoticed.

How will the players play? Modify the model if you deem it necessary.

CHAPTER 13

France 1654

1654, Blaise Pascal and Pierre de Fermat exchanged a series of letters. At that time Fermat was 53 years old and a judge in Toulouse. He did mathematics in his spare time, with extraordinary success. A former child prodigy, Pascal was 31 years old and living in Paris. In the letters they discuss a problem posed to Pascal by the Chevalier de Méré, a frequent gambler. Pascal made an attempt at a solution and wrote it to Fermat, probably for confirmation but maybe also for help. We can only guess, since the letter is lost.

Developing scientific ideas in letters instead of papers was not unusual at that time. Scientific journals did not yet exist, the first ones appearing only in 1665. Books were written to communicate ideas, but to develop new results, communicate them to others, and gain credit for a new discovery, scientists and mathematicians wrote letters to each other. Pascal and Fermat had never met, and addressed each other politely and formally as “Monsieur”.

The following question was discussed in the letters: Two players, A and B, toss a coin repeatedly. Heads counts a point for A and tails a point for B. They decide to proceed until one of them has collected six points, when the winner would win the stake of 64 units. What if the game is interrupted when the score is 4 to 3? How should the stakes be divided?

For simplicity, we discuss a simpler game where the player first reaching three points wins, see Figure 13.1. Pascal reasoned as follows: If the points are even, as 2 : 2, 1 : 1, or 0 : 0, then by symmetry both players should expect the same amount of money, 32. Then Pascal considered the case where player A has two points and player B one, a score of 2 : 1. If A gets the next point, A would win, giving a payoff of 64. If B wins the next round, the score would be 2 : 2, giving an expected value of 32 for each. Since both outcomes have 50% probability, Pascal computed the expected value for A as $0.5 \cdot 64 + 0.5 \cdot 32 = 48$. For a score of 2 : 0,




Figure 13.1. Who wins three rounds wins the game

that Pascal had moved to his birthplace Le château de Bien-Assis close to Clermont-Ferrand (which is about 374 km from Toulouse), Fermat initiated contact in a letter in which he suggested a meeting half way between them. In his response, Pascal said that geometry is “the very best intellectual exercise”, but that now he found it “useless”. He continued: “Although I call it the best craft in the world, it is after all only a craft, and I have often said it is fine to try one’s hand at it but not to devote all one’s power to it. In other words, I would not take two steps for geometry...” [PF1654] That settled the question of a possible meeting.

Exercises

1. In the coin tossing game, where the player reaching first seven heads or tails wins, how should the stake of 128 be divided if the game is interrupted with the score 6 : 5? What if is interrupted at 6 : 4? What if it is interrupted at 6 : 3? Use the first approach discussed.
2. In the coin tossing game, where the player reaching first seven heads or tails wins, how should the stake of 128 be divided if the game is interrupted with the score 4 : 2? Use Pascal’s triangle.
3. A game is played where a die is tossed repeatedly. If the die shows a 2 or less, player A gets a point, otherwise player B gets a point. The player first reaching seven points wins. How should the stake of 100 be divided if the game is interrupted with the score 6 : 5? What if is interrupted at 6 : 4? What if it is interrupted at 6 : 3? Use the first approach discussed. Why can’t Pascal’s triangle be used?
4. A game between three players is played where a die is tossed repeatedly. If the die shows a 2 or less, player A gets a point, if it shows a 3 or a 4, player B gets a point, otherwise player C gets a point. The player first reaching five points wins. How should the stake of 100 be divided if the game is interrupted with the score 4 : 4 : 3? What if is interrupted at 4 : 3 : 3? What if it is interrupted at 4 : 3 : 2?

CHAPTER 14

Example: DMA Soccer I

Question: Is this the mathematician talking?

Slomka: I hate mathematics.

Question: But you majored in mathematics?

Slomka: Yes, since I didn't know in advance that it is so cruel.

Interview with Mirko Slomka, coach of Schalke 04,
in the newspaper *Tagesspiegel*, 2007.

Prerequisites: Chapters 1, 2, and 12.

Imagine you are in a poker tournament and the organizer approaches your chair, taps you on the shoulder, and tells you to pass less often. Would you be amused? In a liberal society, we have become accustomed to the fact that organizations and governments leave us freedom. Of course, there are always rules such as pay your taxes, don't speed, and don't smoke during class, but we want them to be minimal, and we enjoy the freedom of selecting our own moves. In many games, players' behavior can be influenced in subtle ways by those in charge. All they need to do is to change the rules. Increase taxes for flying and the people will more often travel by car or train.

In 1994 FIFA, the World Soccer Association, changed the rules for assigning points in soccer tournaments. Before that, a win counted as 2 points, a draw 1 point, and a loss as 0 points. With the new rule, a win counted 3 points. FIFA's aim was to encourage teams to play more offense, after years of defensive, unattractive play. In this chapter we will discuss a game—not soccer but still a game that has defense, midfield, and attack—where changing the point system will influence how coaches set up their teams.

The game **DMA Soccer** (for defense, midfield, and attack) is played on a field divided into three subfields called left, midfield, and right. Two teams, A and B, play against each other. Each team has five players, and each player goes into a subfield and stays there during the whole game. (The requirement is weakened in Chapter 34, where we allow some movements between fields.) At the beginning of the game, and after each goal, the referee throws the ball into the midfield. One player in the midfield will catch the ball. In DMA soccer we always assume that if the ball is thrown into a subfield, then one player there catches it.

- If an A-player in the midfield has the ball, he or she throws it into the right field, and if an A-player catches it there, it counts as a goal for A. If a B-player catches it in the right field, the ball is thrown back into the midfield.

- In the same way, if a B-player in the midfield has the ball, he or she throws it into the left field, where it counts a goal for B if a B-player catches it there, or if an A- player catches it there, it is returned to the midfield.

A ball may never be thrown directly from left to right field, or from right to left field. Starting from when the ball is in the midfield, the time until the ball returns to the midfield or a goal is shot is called a *round*.

What is the best strategy for a team? The coach should select players who can catch and throw the ball well, and who are alert, fit, and have good eyes. But what exactly does this mean? And should different types of players be selected? Even more difficult is to tell how the individual players should play, where they should stand, how they should move, and where to throw the ball. We are not saying how to play the game. We consider only how the coach should distribute the players into the subfields, assuming that the players have more or less equal abilities. We model the game as a simultaneous one between the two coaches, each coach having only one move—choosing the distribution of players into the parts of the field. You are not allowed to react to the distribution of your opponent's team. You have to announce your distribution before the game, without knowing the other team's distribution.

Which distribution do you think is best?

If you want to simulate the game, try the applet [DMAstatic](#).

14.1 1-Round 2-Step Experiment for Given Player Distributions

Each distribution of the five A-players is described by three numbers: DA (for defense) is the number of A-players in the left field, MA (for midfield) is the number of A-players in the midfield, and AA (for attack) is the number of A-players in the right field, and $DA + MA + AA = 5$. Also, DB , MB , and AB are numbers of B-players in the right, middle, and left field, and $DB + MB + AB = 5$. So, each distribution can even be specified by a pair of numbers, (DA, MA) for the A-player, and (DB, MB) for the B-player. There are 21 pairs of numbers with sum at most 5.

If team A uses distribution (DA, MA) and team B uses distribution (DB, MB) , how likely is a goal for A and how likely is a goal for B in a round? A round is defined by two steps: the ball moving out of the midfield, and then moving to the goal or back to the midfield. We make the assumption that all players are equal: if a ball is thrown into a subfield, each of the players there has an equal probability of catching it. Then the probability that the goal initially goes to the left field is $MB/(MA + MB)$. If this happens, the probability of a B-player catching the ball in the left field, and therefore for a goal for B is $AB/(AB + DA)$, where $AB = 5 - DB - MB$, and $AA = 5 - DA - MA$. See the probability tree in Figure 14.1

The probability for a goal for B is

$$p_B = \frac{MB}{MA + MB} \cdot \frac{AB}{AB + DA} = \frac{MB}{MA + MB} \cdot \frac{5 - DB - MB}{5 - DB - MB + DA}.$$

In the same way, the probability for a goal for A equals

$$p_A = \frac{MA}{MA + MB} \cdot \frac{AA}{AA + DB} = \frac{MA}{MA + MB} \cdot \frac{5 - DA - MA}{5 - DA - MA + DB}.$$

There could also be no goal in the round. This corresponds to the second and third outcome in the tree, and its probability is

$$p_0 = 1 - p_A - p_B.$$

In cases where one field has no players, the rules dictate that the ball is moved with 50% probability for each option. Either to the left or right field, if the midfield is void (i.e., if $MA = MB = 0$), or the choices




Figure 14.1. The probability tree

are goal or midfield (if $AA = DB = 0$ or $AB = DA = 0$). Thus if somewhere in the formulas division by 0 would occur, the fraction is replaced by $1/2$.

Let us compute p_A and p_B for distribution $(1, 3)$ against distribution $(1, 2)$. We get $p_A = \frac{3}{3+2} \cdot \frac{1}{1+1} = 0.3$, and $p_B = \frac{2}{2+3} \cdot \frac{2}{2+1} = 4/15 = 0.267$.

Computing the expected value of goals for team A in one round is simple. We have a value of 1 with probability p_A and a value of 0 with probability $1 - p_A$, so the expected value is $p_A \cdot 1 + (1 - p_A) \cdot 0 = p_A$. In the same way, p_B is the expected value of goals for team B in the round, and the expected goal difference is $p_A - p_B$.

Student Activity Use the applet [DMAstatic](#) to simulate 100 rounds where A plays distribution $(1, 3)$ and B plays distribution $(1, 2)$. How many goals did Team A score in these 100 rounds? How many did Team B score? How many goals were Teams A and B expected to score according to the formulas?

14.2 Expected Goal Difference for the One-Round Game

The goal of soccer is to score goals and prevent the opposing team from scoring goals. Thus the goal difference might be the payoff we want to maximize, although we will see in Section 4 that this approach is not quite the best one. It is straightforward to compute the expected values of the goal difference in one round for different player distributions, as has been done in the Excel sheet [DMA1.xlsx](#) on the sheet “matrix”. If the payoff is the goal difference, then we have a zero-sum simultaneous game, where the payoffs are determined by random experiments based on the choices of the teams. The best response digraph, actually the condensed version used for symmetric games, is in Figure 14.2. When an arrow goes from vertex $(1, 1)$ to vertex $(0, 3)$, then $(0, 3)$ is Ann’s best response to Beth’s move $(1, 1)$, but $(0, 3)$ is also Beth’s best response to Ann’s move $(1, 1)$. The 21 moves for each player are in a grid, where the three corners of the triangular structure correspond to distributions with strong attack, strong midfield, and strong defense, and where the relation of a distribution to the corners indicates its emphasis on defense, midfield, or attack.

From the digraph it follows that the matrix has four pure Nash equilibria, the first $(1, 3)$ versus $(1, 3)$, the second $(2, 2)$ versus $(2, 2)$, and the third and fourth $(1, 3)$ versus $(2, 2)$ and $(2, 2)$ versus $(1, 3)$.

There is weak but no strong domination. Note that $(2, 2)$, which is part of a pure Nash equilibrium, is weakly dominated by $(1, 3)$. After deleting weakly dominated moves, the moves $(0, 1)$, $(0, 2)$, $(0, 3)$, $(0, 4)$, $(1, 1)$, $(1, 2)$, and $(1, 3)$ remain for both players. In this 7×7 matrix, the moves $(0, 1)$, $(0, 2)$, and $(1, 1)$ are weakly dominated. After deleting them as well, the resulting 4×4 matrix has a weakly dominating move




Figure 14.2. Best response digraph for the goal difference

(1, 3). Therefore the IEWD equilibrium is the move (1, 3).

There are two maximin moves for both players: (1, 3) and (2, 2). The maximin analysis is done on the Excel sheet [DMA1.xlsx](#) mentioned on the “matrix” sheet. In column Z the minima of the row values are computed, and their maximum, the maximin value, is in cell Z9. In row 31 the maxima of the column values are computed for Beth. Since this is a zero-sum game so far, Beth’s payoffs are the negatives of Ann’s payoffs. Therefore the minima of Beth’s values are the maxima of the corresponding values for Ann. The minimum of the maxima values is in cell D31. The red values show where the maxima or minima are attained.

It is not surprising that the midfield is important, but it is curious that attack and defense are not symmetric. The move (1, 2), which is the symmetric counterpart to (2, 2), never occurs as a recommendation for a move.

14.3 3-Rounds Experiment for Given Player Distributions

We move back from games to experiments. Again we assume in this section that both player distributions are already fixed. Now we look at a three-round experiment consisting of three rounds of the basic experiment discussed in Section 14.1. The outcome is a result of the game, as 1-1, 1-2, or 1-0. Since no more than three goals can be achieved in three rounds, there are ten outcomes, as shown in the probability digraph in Figure 14.3. There are three kinds of arcs, those of probability p_A , which are colored green, those of probability p_B , which are colored red, and those of probability p_0 , which are colored black. Recall that the numbers p_A, p_B, p_0 are the probabilities for a goal for team A, for a goal for team B, and for no goal. Formulas for p_A, p_B, p_0 were derived in Section 14.1.

Let us compute the probabilities of getting the outcomes. We want to assign probabilities to all situations in the probability digraph. Let us abbreviate having a score of a to b after n rounds by $a:b|n$, and let $p(a:b|n)$ denote the probability with which it will occur. The probability of $0:0|0$ is 1 and $1:0|1$, $0:0|1$, and $0:1|1$ occur with probabilities p_A, p_0 , and p_B . More interesting are the scores after two rounds.

- There is only one way to obtain $2:0|2$ —from $1:0|1$ by scoring a goal for A. This is clear since there is only one directed path from the start to $2:0|2$. Therefore $p(2:0|2) = p(1:0|1) \cdot p_A = p_A \cdot p_A$.
- On the other hand, $1:0|2$ can be obtained in two ways: from $0:0|1$ and getting a goal for A, or coming from $1:0|1$ and getting no goal. Thus $p(1:0|2) = p(0:0|1) \cdot p_A + p(1:0|1) \cdot p_0 = 2 \cdot p_A \cdot p_0$.
- In the same way we obtain $p(1:1|2) = p(1:0|1) \cdot p_B + p(0:1|1) \cdot p_A = p_A \cdot p_B + p_B \cdot p_A = 2 \cdot p_A \cdot p_B$,
- $p(0:0|2) = p(0:0|1) \cdot p_0 = p_0 \cdot p_0$,
- $p(0:1|2) = p(0:0|1) \cdot p_B + p(0:1|1) \cdot p_0 = p_0 \cdot p_B + p_B \cdot p_0 = 2 \cdot p_0 \cdot p_B$,




Figure 14.3. The probability tree for three-round DMA soccer

- $p(0:2|2) = p(a:1|1) \cdot p_B = p_B \cdot p_B.$

In the same way, the probabilities for the outcomes after three rounds can be obtained:

- $p(3:0|3) = p(2:0|2) \cdot p_A = p_A^3,$
- $p(2:0|3) = p(2:0|2) \cdot p_0 + p(1:0|2) \cdot p_A = 3 \cdot p_A^2 \cdot p_0,$
- $p(2:1|3) = p(2:0|2) \cdot p_B + p(1:1|2) \cdot p_A = 3 \cdot p_A^2 \cdot p_B,$
- $p(1:0|3) = p(1:0|2) \cdot p_0 + p(0:0|2) \cdot p_A = 3 \cdot p_A \cdot p_0^2,$
- $p(1:1|3) = p(1:0|2) \cdot p_B + p(1:1|2) \cdot p_0 + p(0:1|2) \cdot p_A = 6 \cdot p_A \cdot p_B \cdot p_0,$
- $p(0:0|3) = p(0:0|2) \cdot p_0 = p_0^3,$
- $p(0:1|3) = p(0:0|2) \cdot p_B + p(0:1|2) \cdot p_0 = 3 \cdot p_B \cdot p_0^2,$
- $p(1:2|3) = p(1:1|2) \cdot p_B + p(0:2|2) \cdot p_A = 3 \cdot p_A \cdot p_B^2,$
- $p(0:2|3) = p(0:1|2) \cdot p_B + p(0:2|2) \cdot p_0 = 3 \cdot p_B^2 \cdot p_0,$
- $p(0:3|3) = p(0:2|2) \cdot p_B = p_B^3.$

The probabilities could be derived more elegantly using trinomial coefficients.

Since the first four outcomes are a win for A, the next two a draw, and the last four a win for B, the probabilities for each of these events can be found. For instance, A wins with probability

$$p_A^3 + 3p_A^2p_0 + 3p_A^2p_B + 3p_Ap_0^2.$$

The probability of a draw is

$$6p_Ap_Bp_0 + p_0^3.$$

The probability of a loss for Ann is

$$3p_Bp_0^2 + 3p_Ap_B^2 + 3p_B^2p_0 + p_B^3.$$

Student Activity Simulate ten games in the three-round version, using the applet [DMAstatic](#), where A plays distribution (1, 3) and B plays distribution (1, 2). How often did Team A win? How often did Team B win? Compare the values with the theoretical values just obtained.

As in Section 14.1 we could ask about the expected numbers of goals for A and B. If you repeat an experiment n times, the expected sum of the random variable in the n experiments is n times the expected value of the random variable in a single experiment. Therefore the expected number of goals for A in the three-round experiment is $3p_A$, and the expected number of goals for B is $3p_B$. The expected goal difference is $3(p_A - p_B)$.

The payoff for a soccer match is not the number of goals, or the goal difference, but the points you get for win, draw, and loss. Usually you get 3 points for a win, 1 for a draw, and no points for a loss. Therefore soccer is not a zero-sum game! The expected value of points for Ann is

$$[3p_A p_0^2 + 3p_A^2 p_0 + p_A^3 + 3p_A^2 p_B] \cdot 3 + [p_0^3 + 6p_A p_B p_0] \cdot 1.$$

and the expected value of points for Beth is

$$[3p_B p_0^2 + 3p_B^2 p_0 + p_B^3 + 3p_B^2 p_A] \cdot 3 + [p_0^3 + 6p_B p_A p_0] \cdot 1.$$

For an example, consider Ann's distribution (1, 3) versus Beth's distribution (1, 2). We have seen in Section 14.1 that $p_A = \frac{3}{10} = \frac{9}{30}$, $p_B = \frac{4}{15} = \frac{8}{30}$, and therefore $p_0 = \frac{13}{30}$. Then the probability of A winning is

$$p_A^3 + 3p_A^2 p_0 + 3p_A^2 p_B + 3p_A p_0^2 = \frac{9^3 + 3 \cdot 9^2 \cdot 13 + 3 \cdot 9^2 \cdot 8 + 3 \cdot 9 \cdot 13^2}{27000} = 0.385,$$

the probability of a draw is

$$6p_A p_B p_0 + p_0^3 = \frac{6 \cdot 9 \cdot 8 \cdot 13 + 13^3}{27000} \approx 0.289,$$

and the probability of Beth winning is

$$3p_B p_0^2 + 3p_A p_B^2 + 3p_B^2 p_0 + p_B^3 = \frac{3 \cdot 8 \cdot 13^2 + 3 \cdot 9 \cdot 8^2 + 3 \cdot 8^2 \cdot 13 + 8^3}{27000} \approx 0.326.$$

The expected value of points Ann wins in this game is $3 \cdot 0.385 + 1 \cdot 0.289 \approx 1.44$, and the expected value of points Beth wins in this game is $3 \cdot 0.326 + 1 \cdot 0.289 \approx 1.27$.

14.4 Static Three-round Game

Let us again turn the focus from probability experiments, where we cannot act but only observe, towards games. The formulas for the expected number of points derived in the previous section have to be evaluated for all possible values of DA, MA versus all possible values of DB and MB. This is a task best left for Excel. In the file [DMA2.xlsx](#) on the sheet "3-round", the expected number of points are computed. You can change the number of points assigned for win and draw. The best responses are computed on the sheet "BestR".

Using the file, let us compare three different rules. The old 2-1 rule, assigning 2 points for a win and 1 point for a draw, the new 3-1 rule, assigning 3 points for a win and 1 point for a draw, and a 3-2 rule. It turns out that the differences in the best response digraph are small. The best response digraph for the 2-1 rule looks exactly like the one considered above for the goal difference. The two other best response digraphs are shown in Figures 14.4 and 14.5.

Let me describe the differences. Going from the 2-1 to the 3-1 rule, the best response to (0, 3) is no longer (1, 3) but (0, 4), (2, 2) is no longer a best response to (1, 3), and (1, 3) and (2, 2) are no longer best responses to (2, 2). Going from the 2-1 to the 3-2 rule, (1, 3) is no longer a best response to (1, 3), and (1, 2) and (1, 3) are no longer best responses to (2, 2). Under the 3-1 rule adopted in 1994, the distribution (2, 2) loses some of its attractiveness, and the only pure Nash equilibrium is (1, 3) versus (1, 3). Under the 3-2 rule, the only pure Nash equilibrium is (2, 2) versus (2, 2).

In our simple model, the shift from the 2-1 rule to the 3-1 rule discourages teams from playing the more defensive distribution (2, 2) and encourages the more offensive distribution (1, 3).




Figure 14.5. Best response digraph for 3 rounds with 3-2 rule

14.5 Static Nine-round DMA Soccer

How would the analysis of the static nine-round DMA Soccer be more complicated? Instead of having 10 suboutcomes, we would have 55: 0-0, 1-0, 2-0, 3-0, 4-0, ..., 0-8, 1-8, 0-9. The formulas for the probabilities attached to them would not look that much different to the three-round case. A win for Ann can be obtained in 25 ways, and a draw in 5 ways, so the probability of a win for Ann would be the sum of 25 terms, and for a draw the sum of 5 terms. This is not totally out of reach.

Exercises

1. Let team A have three defense, one midfield, and one attack players, and team B have two defense, two midfield, and one attack players.
 - a) What is the probability of a goal for team A in a round?
 - b) What is the probability of a goal for team B in a round?
 - c) What is the expected number of goals for team A in a three-round play?
 - d) What is the probability of a result of 2-1 for team A in a three-round play?
 - e) What is the probability of a result of 1-1 in a three-round play?
 - f) What is the probability of a draw in a three-round play?
 - g) What is the expected number of points won for team A in a three-round play?
2. Assume each team has six players, and A has three defense, two midfield, and one attack players, and Beth has two defense, two midfield, and two attack players.
 - a) What is the probability of a goal for team A in a round?
 - b) What is the probability of a goal for team B in a round?
 - c) What is the expected number of goals for team A in a four-round play?
 - d) What is the probability of no goal in a round?
 - e) What is the probability of a result of 2-1 for team A in a four-round play?
 - f) What is the probability of a result of 1-1 in a four-round play?
 - g) What is the probability of a draw in a four-round play?
3. We have seen that best strategies for goal difference look slightly more defensive, more (2, 2)-ish, than for winning under the 3-1 rule. Look at three existing soccer leagues. Which team did finished first, and which team had the largest goal difference? Does this confirm what we have found?

4. Confirm that the expected goal difference for move (3, 1) versus move (2, 2) is zero. Calculate the probability that the player playing (3, 1) wins a three-round game against one playing (2, 2), and the probability that the (3, 1) player loses. Confirm that the probabilities are the same. Why, in the 3 points for a win and one point for a draw scheme, is the best response to (2, 2) only (3, 1)?

Project 22

DMA6* SOCCER In this variant, each team has six players, and must send at least one player into every subfield. Play the game in [DMA6static](#). Draw the best response digraph and analyze the game for the variant where the payoff for each player is goal difference. This is the same point of view taken in the second section of this chapter. Modify the Excel file [DMA1.xlsx](#). Also try to analyze the three-round static game, either with 2 or 3 points for a win, or analyze both. Here you would use the Excel file [DMA2.xlsx](#).

Project 23

DMA7* SOCCER In this variant, each team has seven players, and must send at least one player into every subfield. Play the game in [DMA7static](#). Draw the best response digraph and analyze the game for the variant where the payoff for each player is goal difference. This is the same point of view taken in the second section of this chapter. Modify the Excel file [DMA1.xlsx](#). Also try to analyze the three-round static game, either with 2 or 3 points for a win, or analyze both. Here you would use the Excel file [DMA2.xlsx](#).

CHAPTER 15

Example: Dividing A Few Items II

A glimpse into statistics, mechanism design, and incomplete information

Prerequisites: Chapters 1, 8, 12, and 9.

You have two lovely daughters, Ann and Beth. A friend gave you a few presents for them. You don't know how much the presents mean to your daughters, but you want to be fair, and you want them to be as happy as possible. How would you distribute them?

We don't have a game yet, only the description of a problem.

The most important assumption, used in Section 15.4, is that the daughters are so close to each other that they can estimate the value of each item to their sister. The more complicated incomplete information situation where this is not the case is briefly discussed in Section 15.5.

You could distribute the items yourself. Since you don't know about the values of the items to your daughters, it may seem fairest to give about the same number of items to each. However, the outcome is not fair if one daughter is much happier with your choice than the other. Another disadvantage of the method is that it is not efficient—the items are distributed without considering their values to Ann and Beth. Since you don't know these values, but Ann and Beth do, they should be involved in the distribution process. This will make it more likely that they get items that are valuable to them.

You could ask Ann and Beth which items they prefer. But would they tell you the truth, or would they lie tactically? Instead of asking them, you will set up a procedure, a game, to select the items. Your task here is not to play the game successfully—this is the task for Ann and Beth. Your task is to design a game that is good in a sense that will be explained in Section 15.1.

We will present three families of simple games. Then we will compare them for seven items, using computer simulations.

15.1 Goals of Fairness and Efficiency

We assume that the values of the items are on a ratio scale, and that the total utility for a player is the sum of her values of the items she gets. There are no two items that only make sense together, like left and right shoe, and no item lowers the worth of another item. Both assumptions are strong.

15.1.1 Fairness

Look at the following example:

values	D	E	F	G	H
Ann	12	62	42	54	30
Beth	51	87	54	39	69

In this example, Ann doesn’t seem to value the items much. For each except item G, Beth has a higher value than Ann. Should we maximize the sum of both their values and give all items except item G to the more grateful Beth? Or, for fairness, should we give two or more items to the more reserved Ann?

The answer depends on whether the values are absolute or subjective. In the first case, which we call interpersonal comparability of utilities, yes, giving all items except G to Beth would be probably wise. In the second case, we may want to make the sum of values the same for both players. We can use **relative values**, the ratios of individual values to the sum of all values. This is the approach we take.

In the example, the sum of the values of the items for Ann is 200, and the sum of the values of the items for Beth is 300. Therefore the relative values are:

relative values	D	E	F	G	H
Ann	$\frac{12}{200} = 6\%$	$\frac{62}{200} = 31\%$	$\frac{42}{200} = 21\%$	$\frac{54}{200} = 27\%$	$\frac{30}{200} = 15\%$
Beth	$\frac{51}{300} = 17\%$	$\frac{87}{300} = 29\%$	$\frac{54}{300} = 18\%$	$\frac{39}{300} = 13\%$	$\frac{69}{300} = 23\%$

Let us call the sum of Ann’s relative values of the items she gets **Ann’s satisfaction**, and the sum of Beth’s relative values of the items she gets **Beth’s satisfaction**. An outcome is fairest if both players achieve the same satisfaction. We call the absolute value of the difference between the two satisfactions the **inequity**. A game is fair if it produces a low average inequity.

15.1.2 Efficiency

Distributions where Ann and Beth get about the same numbers of items, but where Ann and Beth both get the items they prefer least, would obviously be inefficient: By exchanging their items, both players would increase their satisfaction. We call an outcome **Pareto-dominated** if there is another distribution giving a higher satisfaction to both. Our first efficiency-related requirement is that the outcome of the game should not lead to a Pareto-dominated outcome.

We could also look at the sum of Ann’s and Beth’s satisfaction. But this value is still not meaningful enough, since it depends on the data—the values the players attach to the items. If their preferences are close, we may not be able to get the sum of Ann’s and Beth’s satisfaction much higher than 100%. So we rather want to compare the sum of satisfactions to what may be possible.

The **inefficiency** of a distribution of the items to Ann and Beth is the difference between the maximum sum of Ann’s and Beth’s satisfactions and the sum of Ann and Beth’s satisfaction in the distribution. The maximum sum of satisfactions is easy to compute for somebody who, like Santa Claus, knows the values Ann and Beth attach to the items: Each item goes to the player who attaches the higher relative value to it.

In the example, Santa Claus gives items E, F, and G, with relative values 31%, 21%, and 27% to Ann, and Beth gets items D and H with relative values 17% and 23%. Ann’s satisfactions is 79%, and Beth’s is 40%. The sum of the satisfactions is 119%. This is the “Santa Claus” value against which satisfaction sums from other distributions will be compared.

15.1.3 Three Additional Features

We will concentrate on inequity and inefficiency, but in the literature there are other features measuring fairness and efficiency of an outcome.

- An outcome is called **envy-free** if each player prefers what she gets over what the other player gets. This means that the satisfaction for each player exceeds 50%.

- An outcome should not be Pareto-dominated.
- We could also aim to get a distribution where the satisfaction of both is as high as possible, though not necessarily equal. Thus the minimum of Ann's and Beth's satisfaction could be looked at. We call this value **Min Sat**

15.1.4 Mechanism Design

Although the goals are to some extent conflicting, you, the parent, may now have a clearer picture of what you want. Let's assume you want to keep inefficiency and inequity small, in a balanced way. The problem is how to achieve a distribution that does this. You ask your two older brothers for advice. Each has two daughters.

Your oldest brother Adam tells you that he lets his daughters select the items one by one. The oldest daughter has the first choice, and then they alternate until everything is chosen. As little parents' involvement as possible, is his motto. In this way, he declares, he gets efficient outcomes, as if directed by an invisible hand.

Your younger brother Karl tells you that he doesn't believe in invisible hands but only in his own. He always divides the items himself. He admits that he doesn't know the preferences of his children, but insists that this way the outcome is at least fair, fairer than Adam's procedure, which favors the oldest daughter. But later, Adam tells you that Karl's children, unsatisfied with the outcome, have voted for introducing Adam's procedure.

So what do you do? Choose the procedure (or mechanism) Adam has used successfully for so long? Or rather try Karl's mechanism, even though his children didn't like it? As in real life, there are more than two options. We can invent our own mechanism and use it. Before selecting a mechanism, we might test how good several mechanisms are, and then pick the best. This is what we will do in Section 15.4.

Mechanism design is a new branch of game theory, that tries to achieve what we described: Without knowing the preferences of the players, the task is to design a game that achieves some desirable outcomes as closely as possible.

15.2 Some Games

15.2.1 Selecting one by one Games

Here is Adam's procedure to let the two players choose items one by one.

Starting with Ann, both alternate in selecting one item until all are chosen. If there are five items, we call the game ABABA, for six items ABABAB, and so on.

Since Ann starts, she will have some advantage, which is larger for an odd number of items, since then Ann gets more items than Beth. For that reason, we consider variants of these games, where at some stages players are allowed to choose more than one item. For instance, ABBABA is a game with six items where first Ann chooses one item, then Beth chooses two, then Ann chooses one, then Beth one, and Ann takes the last. In the same way the games ABBAAB could be defined for six items, and the games ABBAB, ABBAA, and ABABB for five items. Each of these games is a sequential game with perfect information. They have been analyzed in Chapter 9.

15.2.2 Cut and Choose

This game is a well-known method for dividing something that can be divided continuously, like a pizza or a cake. Ann cuts the item into two parts, and Beth chooses whatever part she wants. Our discrete version is as follows:

CUT AND CHOOSE C&C Ann divides the items into two heaps. Beth decides which heap she wants, and Ann gets the other.

This is a sequential game with two moves, perfect information, and no randomness. For n items, there are 2^n possible heaps, and $2^n/2 = 2^{n-1}$ ways of dividing the items into two heaps. So Ann has many possible moves, 64 for $n = 7$, but Beth has only two moves.

In the continuous version, the outcome is envy-free, since everybody gets at least 50% satisfaction. This is not true for the discrete version. If Ann cannot divide the items into two heaps of about the same value for her, she can get less than 50%. On the other hand, Beth never gets less than 50% satisfaction. Though this looks unfair to Ann, in Section 15.4 we will see that the game favors Ann under some assumptions on the distribution of the values to the items.

15.2.3 Random and Exchange

Exchange games start with some distribution, in our case a random one. Then there are a few rounds of exchanges that the players negotiate. There is a variety of rules for these exchanges. We consider this one:

R&E2 The items are distributed randomly, with each player getting about the same number. Then Ann tells Beth which one of Beth’s items she would like to have. Beth tells Ann which one of Ann’s items she would like to have in exchange. If both agree with the deal, the two items are exchanged. Altogether two rounds of negotiations are played, with the player proposing first alternating.

The version with four rounds is called **R&E4**.

15.3 Examples

In this and the next section we assume that Ann and Beth know each other’s preferences. Then the games we described are sequential games with perfect information and can be analyzed using backward induction. We assume that both players are able to do so, and will always play optimally.

Example 1 Look at the following relative values with six items:

relative values	D	E	F	G	H	I
Ann	26%	6%	3%	28%	5%	32%
Beth	21%	16%	23%	14%	8%	18%

Ann essentially values only items D, G, and I. Before we analyze what happens if the games are played with this data, let us first calculate the maximum possible sum of satisfaction. Santa Claus gives each item to the player for which it has the highest relative value. Therefore item D goes to Ann, items E and F to Beth, item G to Ann, item H to Beth, and item I to Ann. Ann’s satisfaction is $26\% + 28\% + 32\% = 86\%$, and Beth’s satisfaction is $16\% + 23\% + 8\% = 47\%$. The total achievable sum of satisfaction is $86\% + 47\% = 133\%$.

Let’s first use the game ABABAB. According to the analysis in Chapter 9, Ann first selects item I, worth 32% to her. Then Beth selects item D with a worth of 21%. Beth does not play greedily at that point. Item F, which is a little more valuable to her, can wait, since Ann does not want it. Then Ann selects item G (28%), Beth selects item E (16%, item F can still wait), Ann selects item H (5%), and

Beth takes the remaining item F (23%). Ann’s satisfaction is 65% and Beth’s satisfaction is 60%. The inefficiency is $133\% - 65\% - 60\% = 8\%$, the inequity is $65\% - 60\% = 5\%$, and it can be seen that the outcome is not Pareto-dominated. This is not a coincidence, because it has been shown [BS1979] that no outcome of a selecting one by one game with both players alternating is Pareto-dominated.

What happens when playing game C&C? By looking into the list of all distributions, one can see that Ann will put items E, G, and I into one heap and items D, F, and H into the other. Since Beth’s relative values of the three items in the first heap add to $16\% + 14\% + 18\% = 48\%$ and in the second heap to $21\% + 23\% + 8\% = 52\%$, she will take the second and obtain a satisfaction of 52%. Ann’s satisfaction is therefore $6\% + 28\% + 32\% = 66\%$. The inefficiency of the outcome is $133\% - 66\% - 52\% = 15\%$, considerably higher than when playing the game ABABAB. The inequity is $66\% - 52\% = 14\%$, also higher than for ABABAB. The outcome is not Pareto-dominated—it never is when C&C is played properly.

The outcome of R&E2 depends on the random move at the beginning, where the items are distributed randomly to Ann and Beth. Assume Ann starts with items E, F, G, and Beth with items D, H, I. Ann is willing to exchange E for anything but H, and to exchange F for anything. But Beth is not much interested in item E (except exchanging with H, which Ann wouldn’t agree to), and will agree only to an exchange of item F for item H. If Ann proposes anything else, Beth will decline, and then propose the exchange F for H herself. Even though Ann might decline this, eventually she will agree. The outcome—E, G, H for Ann and D, F, I for Beth— is not different if we allow more rounds. Ann’s satisfaction is only 39%, Beth’s satisfaction is 62%, and the inefficiency and inequity are 32% and 23%. The outcome is Pareto-dominated by the distribution of items G and I for Ann and items D, E, F, and H for Beth, but this distribution cannot be obtained from a three items versus three items distribution by exchanges of item for item.

Example 2 Here is another example with seven items:

relative values	D	E	F	G	H	I	J
Ann	33%	1%	3%	2%	17%	9%	35%
Beth	11%	18%	21%	19%	10%	8%	13%

Beth is eager to get items E, F, and G. Ann will not mind. The maximum sum of satisfaction is $33\% + 18\% + 21\% + 19\% + 17\% + 9\% + 35\% = 152\%$.

When playing ABABABA, Ann gets items J(35%), D(33%), H(17%), and I(9%) in that order. Beth chooses items F(21%), G(19%), and E(18%) in that order. Ann’s satisfaction is $35\% + 33\% + 17\% + 9\% = 94\%$, and Beth’s satisfaction is $21\% + 19\% + 18\% = 58\%$. The inefficiency is $152\% - 94\% - 58\% = 0\%$, but the inequity is $94\% - 58\% = 36\%$.

Since ABABABA favors Ann, let’s try ABABABB. Ann gets items J(35%), H(17%), and I(9%) in that order. Beth chooses items D(11%), F(21%), G(19%), and E(18%) in that order. Beth chooses item D, that now both want, first. Ann’s satisfaction is $35\% + 17\% + 9\% = 61\%$, and Beth’s satisfaction is $11\% + 21\% + 19\% + 18\% = 69\%$. The inefficiency is $152\% - 61\% - 69\% = 22\%$, but the inequity is low, $69\% - 61\% = 8\%$. The outcome is not Pareto-dominated.

When playing C&C, we get the same efficient outcome as in ABABABA.

Assume that Ann starts with items D, E, and I in R&E2, and Beth with items F, G, H, and J. They will exchange item E for item H, but after that they are stuck. Although Ann wants item J, she has nothing to offer to Beth for it. Ann ends with a satisfaction of $33\% + 17\% + 9\% = 59\%$, whereas Beth gets $18\% + 21\% + 19\% + 13\% = 71\%$. The inefficiency of the outcome is $152\% - 59\% - 71\% = 22\%$, and the inequity $71\% - 59\% = 12\%$. The outcome is Pareto-dominated.

15.4 Comparison of the Games for Seven Items and Complete Information

So which game is the best? This question is ill-posed for two reasons. First, we have to specify what “best” means. It could be any of the criteria discussed, but it could be something else. It is a matter of judgment what society wants, which cannot be resolved mathematically. Second, even if we focus on efficiency, for instance, what game produces the least inefficient outcome depends on the data, on the values Ann and Beth attach to the items.

We have to decide which game to choose while we don’t know Ann’s and Beth’s preferences. That’s the whole point. We want to find out which of the games produces the most efficient outcome for generic preferences, or on average.

Our problem is similar to that of finding what medicine is best for treating patients with a given disease, let’s say high blood pressure (HBP). Different medicines work differently for different patients, but which is best on average? We could pick one patient randomly, test the medicines on him or her, and base the recommendation on the result.

In our problem, we pick a set of preferences randomly. The random preferences are not really picked but created, under an assumption about how they may look. We assume that the value for an item is a random number between 0 and 1, and that all the values are independent.

If we test medicines on one random HBP patient, we get a random result. More meaningful results are obtained by testing more than one patient. The average of the sample would get closer and closer to the average of the whole population of HBP patients as the sample size increases. But we have to make sure that the sample is representative—if all the patients in the sample are children, then the average of the sample may not get close to the average of the whole population. This is done by choosing the patients randomly. One of the most common tasks of statistics is estimating an average in a population by calculating an average in a random sample. We will do this for our example.

I simulated 8000 randomly generated preference patterns for seven items, where the fourteen values of the items for Ann and Beth are independent random numbers with a uniform distribution between 0 and 1. Here are the averages for the features and games, provided both players play optimally.

	Ann’s sat	Beth’s sat	In- efficiency	In- equity	Min Sat	% not Pareto- dominated	% of envy-free outcomes
C&C7	73.3%	54.6%	4.54%	18.77%	54.6%	100%	99.96%
ABABABB	61.6%	65.1%	5.75%	9.57%	58.6%	81.0%	93.9%
ABBABAA	68.4%	58.4%	5.72%	12.47%	57.1%	84.2%	88.5%
ABABBAB	60.2%	66.5%	5.71%	10.53%	58.1%	82.2%	90.7%
R&E2	53.3%	71.3%	7.9%	20.4%	52.1%	75.7%	61.4%
R&E4	54.0%	71.6%	6.9%	20.0%	52.8%	82.7%	63.7%

In addition, for these 8000 data sets the average maximum sum of satisfactions is 132.5%, the average minimum possible inequity is as low as 0.8%, and the average maximum possible minimum satisfaction is 62.4%.

Average inefficiencies and inequities for the games are displayed as the red diamonds in Figure 15.1. C&C seems to be the game promising the least inefficiency, and the turn-taking games, in particular ABABABB, seem to produce the smallest inequity.

How sure can we be that the results will be meaningful not just for our sample of 8000 values but for any random data? Besides the sample size, the **standard deviation** of the sample determines how close the average of the sample is to the average of the population. The formula for the standard deviation looks




Figure 15.1. Average inequality versus inefficiency for some games, $n = 7$

complicated, but the concept is not. The **variance** is the average square of the deviation from the actual number to the mean of the sample. The standard deviation is the square root of the variance. This is not the average deviation between the actual value and mean, which is 0. Excel has a build-in function for standard deviation.

Fact If we find a standard deviation of s for a sample of size n , then with at least 95% probability, the mean of the population does not differ from the mean of the sample by more than $\frac{2s}{\sqrt{n}}$.

The standard deviations for the inefficiency of C&C and ABBABAA are about 0.045 and 0.059. Therefore with 95% probability the inefficiency lies between 4.44% and 4.64% for C&C and between 5.59% and 5.85% for ABBABAA. C&C has almost surely a lower inefficiency than ABBABAA and all the other games considered. The inefficiency of the three turn-taking games cannot be ranked for certain.

The standard deviations for the inequality of the two games with the lowest mean, ABABABB and ABABBAB, are about 0.074 and 0.079. Therefore with 95% probability the inequality for ABABABB lies between 9.4% and 9.74% and for ABABBAB between 10.35% and 10.71% . Thus ABABABB has almost surely a lower inequality than ABABBAB and all the other games considered.

What about other criteria? In Figure 15.2 we graph inequality versus Min Sat. There seems to be a correlation between the variables—the higher the inequality, the lower Min Sat. The relation seems to be almost linear. Consequently, Min Sat seems to measure inequality.

In the same way the percentage of outcomes that are not Pareto-dominated is correlated to inefficiency—the higher the one, the lower the other. The percentage of envy-free outcomes is correlated to inefficiency in the same way. So these two goals both seem to measure inefficiency.

15.4.1 Opposing or Similar Preferences

Do the games perform differently depending on whether Ann’s and Beth’s preferences are similar or opposing? We sorted the 8000 data sets by the correlation coefficient for Ann’s and Beth’s preferences. A correlation coefficient is always between -1 and 1 . It is close to 1 if both preferences are similar, if both play-




Figure 15.2. Min Sat versus inequity, $n = 7$

ers prefer the same items. It is close to -1 if the preferences are opposing, and it is close to 0 if the preferences are independent. It is easier to distribute the items reasonably for opposing preferences. In Figure 15.3 the inequity for four games, C&C, ABABABB, ABBABAA, and R&E4 are displayed for different correlation coefficients, inscribed in bands that show the range where the values would fall with 95% probability. From the graph it follows that ABABABB has lower inequity than the other games almost everywhere, except for very similar preferences.




Figure 15.3. Inequity for preferences from opposing to similar




Figure 15.4. Inefficiency for preferences from opposing to similar

The same is done for inefficiency in Figure 15.4 . Where the 95% certainty bands for inequity were clearly separated, they now overlap. From the graph it follows that for independent or similar preferences (for correlation coefficient of -0.25 or more), C&C has the lowest inefficiency, while for opposing preferences ABABABB and ABBABAA have lowest inefficiency.

15.5 Incomplete Information

What if the players know the values they assign to the items, but not their opponent’s values? Then we have a game of incomplete information.

Even so, a player probably has some belief about the values her sister has. Initial beliefs could change during the play, by inferences made from the decisions of the other player. We will assume that a players believes that the values of her sister are not related to her own values, so the preferences of the players are independent. Other possible beliefs are that the preferences are parallel, or opposed, but they are more difficult to analyze.

With these beliefs of independent preferences, isn’t it rather obvious how players play in one by one games like “ABABABB”? They would play what we called greedy strategy in Chapter 9—choosing the item with the largest value. There is no reason to play tactically and choose something of lesser value, since no prediction is possible about what the other player will choose. All items could be chosen with equal probability.

C&C is more difficult to analyze with incomplete information than with complete information. If Ann naively believes that Beth would take any heap with the same probability, no matter which heaps are offered, then any heap division would do. If Ann’s satisfaction with one heap is x then Ann’s satisfaction with the other heap is $1 - x$. The expected value of satisfaction would be $50\% \cdot x + 50\% \cdot (1 - x) = 50\%$. All Ann knows about Beth’s preferences is that a larger heap is more likely to have a larger value for Beth than a smaller one. Calculating these probabilities is rather difficult, involving calculus, but it can be shown that only with about 0.03% probability a one-item heap has a higher value for Beth than the other 6-item heap, with about 2.4% probability Beth prefers a 2-item heap over the other 5-item heap, and with about 26% probability

Beth prefers a 3-item heap over the other 4-item heap. We assume always that the values of the items have independent uniform distributions.

Assume as an example that Ann’s values for items D, E, F, G, H, I, and J are 7, 5, 2, 2, 1, 1, and 1. If Ann puts D, E, F into one heap (of value 14, satisfaction 14/19) and G, H, I, J into the other (of value 5, satisfaction 5/19), then with 74% probability Beth will choose the 4-item heap, and Ann will get the 3-item heap, valuable for her. Ann’s expected satisfaction is $74\% \cdot 14/19 + 26\% \cdot 5/19 \approx 61.4\%$. If Ann however puts D and E into one heap (of satisfaction 12/19), and the other 5 items into the other (of satisfaction 7/19), then Ann will almost surely get the valuable 2-item heap, and her expected satisfaction is $97.6\% \cdot 12/19 + 2.4\% \cdot 7/19 \approx 62.5\%$, which is even better. Creating a 1-item heap of item D makes no sense for Ann, however, since her satisfaction with that heap would be only 7/19, smaller than 50%.

Analyzing random and exchange games in case of incomplete information is too complicated. After every exchange, the beliefs of the players will change. Since they were not successful for complete information, we will skip them.

I also looked at the outcomes for the remaining games for the 8000 random data sets. The outcomes with respect to inefficiency and inequity can be seen in Figure 15.1, except for C&C, which has an average inefficiency of 14.1% and is not in the window for this reason. Among the selecting one by one games, ABABBAB seems to be best, being even more efficient than the complete information versions. Thus in the complete information case, it would be beneficial for both players to stick to the greedy strategy. However, as with the Prisoner’s Dilemma—they would not stick to it. The result could be expressed by saying that complete information about the other players may sometimes not be beneficial to the players.

Exercises

1. How likely is the sum of three random numbers to be larger than the sum of four random numbers? By random numbers we mean numbers between 0 and 1 with each number equally likely. Simulate it using an Excel sheet. The Excel function “=Random()” creates random numbers. Type it into seven cells, and sum the first three and the last four. Whenever you change something on the sheet, new random numbers are chosen. Count how often the first sum is larger than the second.
2. We are given:

relative values	D	E	F	G	H	I
Ann	10%	10%	30%	20%	20%	10%
Beth	30%	10%	10%	10%	20%	20%

- Look at the distribution where Ann gets items D, E, F, and Beth gets items G, H, I. Calculate inefficiency and inequity. Show that it is Pareto-dominated.
3. With the same data the game R&E2 is played, with an initial random distribution of Ann holding items D, E, F, and Beth holding items G, H, I. What is the outcome of the game, and what are its inefficiency and inequity? Is the outcome Pareto-dominated?
4. We are given:

relative values	D	E	F	G	H	I
Ann	30%	25%	20%	15%	5%	5%
Beth	15%	20%	5%	25%	25%	10%

Compare inefficiency and inequity for the outcomes obtained by playing the games ABABAB and C&C.

Project 24

Dividing five items A Use the Excel sheet [Dividing5.xlsx](#) to find data where each one of the games C&C, ABABA, ABABB, ABBAB, and ABBAA is more efficient than the others. Do the same for inequity and the Max MinSat parameters. Is it possible to describe the data that are good for the games?

Project 25

Dividing five items B Use the Excel sheet [Dividing5.xlsx](#) to collect data for 100 random preferences. Calculate the averages for the features and graph inefficiency versus inequity, and maybe a few others. What is the conclusion?

CHAPTER 16

Theory 5: Sequential Games with Randomness

RANDOM NIM(n, p) n stones lie on a board. As in NIM, Ann and Beth alternate in removing either one or two stones. The player who has to move but cannot (since there is no stone left) loses, and the payoffs are identical to those of NIM. However between moves, 0 or 1 stones are removed randomly, with probability p and $1 - p$.

Student Activity Play RANDOM NIM(5, 0.5) at least 20 times against the computer in the applet [Nim7Rc](#). Put in the values 5 and 0.5 into the text field before you start. Try to win.

The game is sequential, with two players, but between moves of the players there are the random removals. It is a game that is sequential with randomness, and we discuss them in this chapter.

16.1 Extensive Form Extended

To describe and discuss sequential games with some random moves, we will merge the concept of extensive forms of sequential games as described in Chapter 8 with the concept of probability trees discussed in Chapter 12. In addition to the vertices that correspond to positions of the game where a player makes a decision, we need **random vertices**. There the random moves are performed and there are arcs to other vertices from them. These are the positions that can be reached from these random positions. The arcs are labeled by the appropriate probabilities to reach the other positions from the random position. The probabilities attached to the outgoing arcs must sum to 1. The probabilities are known to every player; this is part of the complete information requirement. Thus we have three kinds of vertices: random vertices, position vertices belonging to a player, and end vertices. The start vertex is either a random vertex or belongs to a player.

The digraph for RANDOM NIM(5, 0.5) is in Figure 16.1. The random vertices are in gray and the probability labels of the random moves are suppressed, since they all have the value 0.5. The number shown at a vertex indicates the number of stones still on the board.

16.2 Analyzing the Game: Backward Induction again

We will analyze such games using backward induction. At each vertex we write the *expected* payoffs of the players. Recall that expected payoffs are meaningful only if we have at least an interval scale for them. How should we attach payoffs at random vertices? If all successors of such a vertex have assigned payoffs, we would just use the expected values.




Figure 16.1. Extensive Form of RANDOM NIM(5, 0.5)

Procedure for Backward Induction with Randomness The goal is to have expected values for each player's payoff assigned at every vertex. As long as not all vertices have values attached, do the following:

- (1) Find a vertex with no values attached, but whose successors have values attached. Until all vertices have values, there is one. Call it V .
- (2) At V , either one player, say X has to move, or it is a random vertex. In the first case do step (P3), in the second case step (R3).

(P3) If V is a player vertex belonging to player X , then we proceed as in sequential games without randomness. We identify the successor vertex of V for which the value for player X is largest. This is the vertex where player X wants to move (since the values will turn into payoffs eventually). Call the other vertex W . Do step (P4).

(P4) Copy the values of vertex W to vertex V . Go back to step (1).

(R3) If V is a random vertex, then successors may occur with the given probabilities. Therefore the value at V for a player is the expected value of the values of that player at the successor vertices, i.e., the sum of the products of probabilities to arrive there and the values. Compute the expected value. Go back to step (1)

Such games have been systematically investigated by Kuhn. Zermelo's result holds for them, just for expected payoffs.

Theorem [K1953] Every finite perfect-information sequential game can be analyzed using backward induction. The recommended moves have the property that no player can gain by changing some of her moves, provided other players keep playing their moves.

Although Nash equilibria have not been defined for sequential games yet, this condition has the same flavor.

The expected values and the backward induction choices in our example are shown in Figure 16.2. Ann has a slight advantage.

16.3 Decision Theory: Alone against Nature

In this section we look at 1-player games. They are more optimization problems than games, but backward induction can be demonstrated nicely by them.




Figure 16.2. Extensive form of RANDOM NIM(5, 0.5) after backward induction

Example 1 1-CANDIDATE QUIZ SHOW You play a quiz show with up to six questions. The questions are worth \$1, \$1, \$2, \$4, \$8, and \$16. After each round, you have the option to stop and take the money you have won so far, or hear the next question. If you have heard the next question, there is no way out—you have to answer it correctly to go into the next round. If you give the wrong answer, all your money is lost. Suppose that in round n the probability of giving a wrong answer is $2/(9 - n)$. When do you stop?

Student Activity Play the game in applet [1Candidate](#).

Figure 16.3 presents the extensive form as a game tree. At each round you have the option to attempt the next question or to stop. If you stop, an edge goes to an end position, and the number written on it shows how much you have won. If you attempt the next question, an edge goes to a gray vertex, where a random experiment is performed. It has two outcomes—the answer can be correct or wrong. The probabilities for the outcomes are written below the gray edge. If the answer was wrong, you have again reached an end position, with a payoff of 0. Otherwise, you face the decision whether to stop or to attempt the next question.

We fill in the missing expected payoffs starting at the end. If you have reached the rightmost gray random vertex, i.e., if you have answered the first five questions successfully and attempted the last one, your expected payoff will be $\frac{2}{3} \cdot 0 + \frac{1}{3} \cdot 32 = \frac{32}{3}$. See Figure 16.4. Therefore, at your last decision vertex, you have to choose between an expected payoff of $\frac{32}{3}$ if you attempt the last question, and a sure payoff of 16 if you stop. Since $16 > \frac{32}{3}$, it is rational to stop and the expected payoff is 16. At the second to last random vertex, if you are trying to answer the fifth question, your expected payoff is $\frac{2}{4} \cdot 0 + \frac{2}{4} \cdot 16 = 8$. Therefore if you face the decision whether to allow the fifth question (after answering four questions successfully), you have a choice between an expected payoff of 8 if you attempt the next question, and




Figure 16.3. Game tree of 1-CANDIDATE QUIZ SHOW




Figure 16.4. The right part of the game tree with expected payoffs filled in

a sure payoff of 8 if you stop. We assume that you always go for the higher payoff, expected or sure, so you would be indifferent, but expect 8.

If you now go one step further to the left, the expected payoff of the next (random) vertex is $\frac{2}{5} \cdot 0 + \frac{3}{5} \cdot 8 = \frac{24}{5}$, which is larger than the 4 you would get by stopping. So you would try to answer the fourth question, and never try to answer the last question. Continuing, we get an expected payoff of $\frac{12}{7}$ at the beginning.

Example 2 5 ENVELOPES \$10000 is distributed in five envelopes. One envelope, the bomb, remains empty. In the basic version the player knows how much is put into each envelope, let's say \$1000, \$2000, \$3000, \$4000, and 0. The envelopes are shuffled and in each round you, the player, can choose to get one of the remaining envelopes or keep the money you have collected so far and quit. If you ask for a new envelope, one of the remaining envelopes is selected at random and given to you. If this envelope contains nothing (i.e., is the bomb), you lose all the money you collected in the previous envelopes.

Student Activity Play the game in applet [5Envelopes](#).

Instead of giving an analysis, which is left as a project, let us just a few hints. The game tree is large, but since the situations where you have drawn first \$2000, and then \$3000, or where you have drawn first \$3000, and then \$2000 are identical situations for you, we identify them and move to the digraph. If we label the envelopes containing \$4000, \$3000, \$2000, \$1000, and 0 by A, B, C, D, and E, then the vertices where the player has to decide are those where nothing has been drawn yet, the four cases where one of A, B, C, or D has been drawn, the six cases where two of them have been drawn, the four cases where three of them have been drawn, and maybe the case four have been drawn. The other cases, where envelope E was among those drawn, are end vertices and don't require a decision from the player. Thus $1 + 4 + 6 + 4 + 1 = 16$ vertices belong to the player. For each such vertex, provided the player continues, there is a random vertex. Thus the game digraph would have 16 vertices belonging to the player, 16 random vertices, and the 16 end vertices where the player has drawn envelope E and gains nothing.

Example 3 OH-NO The player rolls a die repeatedly. After each roll, as many dollars as the die shows are added to a heap, which initially is empty. The game ends if

- the player decides that the heap contains enough money and takes it, or
- if a 1 is rolled. Then all the money already collected in the heap is lost. This explains the title of the game.

Student Activity Play the game in applet [OhNo](#).

Analysis: There is only one player. Unfortunately the extensive form is infinite, so backward induction cannot be applied, but we can solve the game using a similar idea. Intuition may tell us that we want to roll the die again provided the heap is small, but may want to stop if the heap has reached some size. The question is which heap size?

The positions for the player are determined by the value of the heap obtained, so we denote them by integers. Each position n has an expected value $v(n)$ for the player—the expected payoff the player would obtain playing optimally from then on. Since in every position the player has the option to stop, obtaining the value of the heap as payoff, $v(n) \geq n$. If the player is at position n and continues, he or she will expect

$$\frac{1}{6} \cdot 0 + \frac{1}{6} \cdot v(n+2) + \frac{1}{6} \cdot v(n+3) + \frac{1}{6} \cdot v(n+4) + \frac{1}{6} \cdot v(n+5) + \frac{1}{6} \cdot v(n+6).$$

Since we started just at position n , and didn't use backward induction, we don't know the expected values at positions $n+2, \dots, n+6$ yet, but we know that they are at least $n+2, \dots, n+6$. Therefore, the player at position n who continues to play expects at least

$$\frac{1}{6} \cdot (n+2) + \frac{1}{6} \cdot (n+3) + \frac{1}{6} \cdot (n+4) + \frac{1}{6} \cdot (n+5) + \frac{1}{6} \cdot (n+6) = \frac{1}{6} \cdot (5n+20).$$

The player will continue playing if $\frac{1}{6} \cdot (5n+20)$ is larger than n , i.e., if $5n+20 > 6n$, or $20 > n$.

We don't know what to do when facing heap sizes of 20 or more, but assuming that the best strategy is to continue for heap sizes below a critical value m and stop after, we can see that 20 is the critical size. Thus the best strategy is to roll until you have \$20. Then you stop. The expected value, however, is smaller than 20 if using the optimal strategy. How large is it?

Exercises

1. What are the expected payoffs and best strategies in RANDOM NIM(5, 1)?
2. What are the expected payoffs and best strategies in RANDOM NIM(5, 0.8)?
3. What are the expected payoffs and best strategies in RANDOM NIM(5, 0.6)?
4. What are the expected payoffs and best strategies in RANDOM NIM(5, 0.4)?
5. What are the expected payoffs and best strategies in RANDOM NIM(5, 0.2)?
6. For which p is RANDOM NIM(5, p) fair, i.e., has the same expected payoff to Ann and Beth?
7. What are the expected payoffs and best strategies in RANDOM NIM(5, 0)?
8. What are the expected payoffs in the game in Figure 16.5, if Ann and Beth play optimally? At the beginning, will Ann choose up or down? When Ann has to decide between hot and cold, which would she choose? When she has to decide between north and south, how would she decide? How would Beth decide between left and right, west and east, and fast and slow?
9. Perform a backward induction for the 3-player game given in Figure 16.6.
10. Analyze the 3×3 version of the following game:

Polyomino REC THE SQUARE with Randomness: Two players alternately put dominos into a square. Between moves, a red square is added randomly, but obeying the polyomino rule, which says that at each state the dominos and random squares must form a polyomino—a connected shape. As in the nonrandom version, the player who cannot move loses.




Figure 16.5. A game tree




Figure 16.6. A game digraph

Play the game in the applet [RecTheSquareRandom3](#). Explain what happens if both play optimally. How much will they expect, if a win counts as +1 and a loss as -1? Use the extensive form in Figure 16.7.

11. **5 PIRATES with random selection** Five pirates have to decide how to distribute 20 gold coins they have found. They do it according to the following rules. In each round they sit in a circle and rotate a bottle of rum. The one to which the bottle points proposes a distribution of the coins, which is voted on, with the proposer voting. The proposal is accepted if more pirates are in favor of it, and in case of a tie. So if there are only two pirates the proposal will always win, and if there are four pirates, the proposer needs only one of the others to vote for the proposal. If the proposal wins, everything is settled. If not, the proposer is thrown into the water, and the next round starts with the remaining pirates. Assume that the pirates value their lives at 40 gold coins, and prefer throwing somebody overboard if everything else is equal (they are playing the hostile variant).

What would you propose if the bottle points at you and there are 2, 3, 4, or 5 pirates left?




Figure 16.7. Game Digraph for Polyomino REC THE SQUARE with Randomness

Project 26

JOB INTERVIEWS You are offered four job interviews in four companies A, B, C, and D. Company A pays \$125,000, but you know that your chance of getting the job is only $1/5$. Company B pays \$100,000 with the chance to be $1/4$. Company C pays \$75,000, and the chance to get the job is $1/3$. Company D pays only \$50,000 but your chance is $1/2$. The job interviews are scheduled sequentially. If you are offered a job, you have to decide immediately whether you accept. If you accept, you cancel the remaining job interviews. If you decline, you cannot come back later.

- Assume company B interviews first, company D second, company A third, and company C last. How would you play? Would you accept an offer company B made? What about company D? How much is the expected salary for your strategy?
- Answer the same questions provided company D interviews first, company A second, company C third, and company B last.
- Assume you can decide which company to visit first, which one second, and so on. How would you order them? Explain.

Project 27

5 ENVELOPES How much would you expect to win? Would you expect to win more or less if the \$10000 was distributed differently in the envelopes, say as \$2500, \$2500, \$2500, \$2500, and 0? What if it were distributed as \$100, \$300, \$2600, \$7000, and 0? What about other distributions? Would you play differently then?

In another version, you can distribute the money into the envelopes (with one of them, the bomb, empty) before starting to play. How would you distribute the money, and how would you play?

In still another version, the distribution of the money in the five envelopes is not known to the player. It is known that exactly one envelope is empty and that the other four contain the \$10000. How would the player's expected value and strategy change?

Project 28

OH-NO or OH-NO6 Write a paper on OH-NO or about OH-NO6 which works in the same way, except that the appearance of a 6 ends the game and deletes all accumulated money. You can play the games in the applets [OhNo](#) and [OhNo6](#).

Among others, you could answer the questions:

- What is the best strategy for playing each of the two games? For OH-NO the best strategy has been given already.
- What is the expected payoff when playing the optimal strategy?
- Consider other strategies, as stopping whenever the heap has reached \$12. What is the expected payoff?
- Can you find the expected payoff of a more creative strategy, as stopping if the heap number is even, and proceeding otherwise?
- You could also perform experiments in the applets and compare the theoretical values for expected pay-offs with the empirical results. For this, use 100 plays for one or several of the strategies mentioned.

Project 29

3 × 4 version of Polyomino REC THE SQUARE with randomness: Play the game in the applet [RecTheSquareRandom34](#), and analyze it. Can the 4 × 4 version, see [RecTheSquareRandom4](#), be analyzed too? If not, describe the difficulty you encounter.

CHAPTER 17

Example: Sequential Quiz Show I

Prerequisites: Chapters 8, 12, and 16.

Pre-Class Activity: Every student should bring a hard multiple-choice question with five choices for answers from another class.

SEQUENTIAL QUIZ SHOW(n, m): Three players, Ann, Beth, and Cindy, have a difficult multiple choice question with five choices. Starting with Ann and continuing cyclically, the player whose move it is can answer or wait. If the player answers correctly, the player gets \$ n . If the answer is incorrect, the player pays \$ m and is out of the game. If the player waits, the quiz master reveals a wrong answer (decreasing the number of choices by one), and the next player moves.

Student Activity Play the game several times, using the questions brought to class.

17.1 Candidates with Little Knowledge

It is easy to play the game well if you know the right answer: give it when it is your turn. In the remainder of this chapter we will show how to play the game if the players don't know the right answer. We assume that the question is so difficult that the players don't know which choice is right, nor can they rule out any of them of being right. So, for the players each of the answers has probability $1/5$. All they can do is guess.

Although the procedure of revealing some wrong answers by the quiz master may recall the Monty Hall paradox, it is not related to it. After an answer has been revealed as wrong, either by a candidate betting on it or by the quiz master when a candidate waits, the probability of being right increases for the remaining answers uniformly. This is different from the Monty Hall paradox. Whenever a candidate faces, for instance, three choices, each could be right with probability $1/3$.

The players move one after the other, so this game is sequential. Assuming that the three candidates don't know the answer and know that the others don't know, the simplest way to describe the game is as follows: First one of the five answers is chosen to be the right one with probability $1/5$. Then Ann moves and has six options: answers A, B, C, D, or E, or waiting. If she tries an answer, the game is over if she is right or she is wrong and is out of the game. If Ann waits, one wrong answer is removed at random with equal probability for the wrong answers. Next Beth would have five options. She could try one of the remaining four answers, or wait, and so on. This description of the game lacks perfect information. The players, Ann, Beth, Cindy, don't know what choice was selected to be right. In Chapter 24 we will see how to formulate imperfect information games in extensive form, but now we describe the game differently, with perfect information.

Although one of the answers is correct when Ann makes her first move, we can describe the game without using this knowledge. Actually it is better to describe it from a spectator's perspective, one who doesn't know which answer is right and doesn't care. During the game, the spectator doesn't even keep track of which answers have been revealed to be false—all that counts for him or her is how many options are left, which players are left, and which player is about to move. This is all that is important for predicting the chances of the players. Even who is about to move follows from knowing how many answer options are left—Ann faces five options, Beth four, Cindy three. The two options question, could be asked of Ann, Beth, or Cindy, depending on who is out of the game. Thus the state with five possible answers implies that Ann moves, and Beth and Cindy are still in the game. With four possible answers, Beth moves. Depending on whether Ann is out of the game, we have two such positions. Cindy is still in the game. At all positions with three possible answers, Cindy is about to move. Ann could have left the game, or Beth, or neither, or both. In the last case, Cindy will win by waiting until only one answer is left. Therefore this position could be expressed as a final vertex, as in Figure 17.1. When two answer options are left, Ann, Beth, or Cindy could have to answer. If Ann has to answer, then four possibilities of Beth, Cindy, neither, or both being out of the game are possible. Again the last case can be formulated as a final vertex. If Beth has to choose between two options, then Ann must be out of the game, therefore there are the two cases of Cindy still in the game or out. The second would be formulated as a final vertex. If Cindy has to choose between two options, then both Ann and Beth must be out of the game, so we have another final vertex. The positions with only one answer option left are of course all final vertices. The moves where the players try an answer are succeeded by positions where it is decided randomly whether the answer was correct or not.

The extensive form of the game is in Figure 17.1.

For SEQUENTIAL QUIZ SHOW(10, 4), using expected values and backward induction, it is easy to deduce that Ann and Beth should both wait, and beginning with Cindy everyone tries any answer. The expected payoff for Ann is 2, for Beth it is $10/3$, and for Cindy it is $2/3$.

In SEQUENTIAL QUIZ SHOW(6, 4) all three players Ann, Beth, and Cindy should wait. Ann would try to answer the question in the second round with two options open. Expected payoffs are 1, 3, and 0 for Ann, Beth, and Cindy.

In the Excel Sheet [SeqQuizShow.xlsx](#), the payoffs for the versions of SEQUENTIAL QUIZ SHOW(n, m) are changed automatically, and the game is solved.

17.1.1 More May be Less

Assume the players have been selected for next day's show. They know that Ann will start, and Beth will move second. The rules specify \$750 for a right answer and a loss of \$400 for a wrong one (i.e., we are playing SEQUENTIAL QUIZ SHOW(750, 400)). At the dinner with candidates and the sponsor of the show, the sponsor surprises her guests by saying that she will raise the money paid for the right answer to \$810. The three candidates should be pleased to hear this, right?

Not really. The rather bright candidates recalculate, and then Ann and Beth ask the donor not to increase the prize. Cindy says dully that she doesn't care much, she would rather have another beer. She has been grumpy the whole evening anyway. How could they decline the donor's generous offer?

If you calculate the expected payoffs in the old and new version, using the Excel sheet, Ann's expectation would be reduced from \$175 to \$137, Beth's from \$375 to \$270, and Cindy's expectation would increase from \$0 to \$3, too small to make Cindy enthusiastic about the change. But still, the change and the positive expectation would cause Cindy to change her strategy. Instead of passing when it is her turn, Cindy would try an answer. This is what causes the loss for Ann and Beth.

It is crucial that we assume the game to be non-cooperative. If the players can cooperate, can discuss the game before it starts, and make binding and enforceable contracts about how to play and make side payments, then Ann and Beth could promise Cindy some money if she would keep her passing strategy. Then all three




Figure 17.1. SEQUENTIAL QUIZ SHOW(10, 4)

would profit from the increased prize. But if the game were cooperative they would play differently anyway. Everybody would pass until Beth, at her second chance, would face only one remaining answer. Then the three would split the prize according to a formula they would have determined. Finding one that is perceived to be fair by all players is part of cooperative game theory.

The graph in Figure 17.2 displays Ann’s, Beth’s, and Cindy’s expected payoffs when playing SEQUENTIAL QUIZ SHOW(n , 4) with n varying between 1 and 22. The sum of the expected payoffs is also shown. There are two cases where strategies change and therefore the sum of the expectations drop, even for increasing prize value. One is for $n = 4$, and the other for $n = 8$, as discussed above. Beth’s position of moving second is always preferable.

17.2 One Candidate Knows More

Assume that at the beginning of the show, when the question and the possible answers have been given, one of the players exclaims: “I know one of the answers is wrong”. What changes?




Figure 17.2. Expected payoffs in SEQUENTIAL QUIZ SHOW(*n*, 4)

We assume that the player is not lying, and that the other players believe her. If she is lying, or if there is doubt, or if she knows one of the answers to be false but doesn't tell the others, things get more complicated. Then we have a game of incomplete information, where not all players have full knowledge about the structure of the game or of the other player's payoffs. We will discuss some special cases later without providing general methods.

At the beginning, the clever candidate has some advantage over the others. Later, when one of the answers has been deleted, the clever candidate may maintain the advantage, provided the deleted answer is not the one the candidate knows to be false. But if it is, then the advantage of the clever candidate is gone. Which situation occurs is decided by luck. With three or two answers left the clever candidate may still have her advantage, or may have lost it. Which situation occurs depends on luck again. We will find out how likely each of the situations is.

The other players cannot learn anything worthwhile from the behavior of the player with the additional information. From whether she passes or attempts an answer, they may learn whether the answer the clever candidate knows to be wrong is still on the list, but that doesn't help them to find the right answer.

Surprisingly, the probabilities for the clever candidate maintaining her advantage when only 4, 3, or 2 answers are left, do not depend on who she is. When an answer is eliminated since the clever candidate tried an answer, then the probabilities are no longer relevant—the player with the advantage has either won or is out of the game. When an answer is eliminated by another player trying this answer, or randomly since some player passed, then the known to be wrong answer is eliminated with probability $1/n$, if n answers are left.

The situation can be described by the probability digraph in Figure 17.3. There are seven relevant situations, denoted as nK or nN with $n = 5, 4, 3, 2$, where n denotes the number of answers left, K that the




Figure 17.3. The probability DAG

clever player still has her advantage, and N that she has lost it, because the answer she knows to be wrong is no longer on the list. State 5N is impossible. Starting with state 5K, the states 4K or 4N could emerge. The transition to 4K has probability $3/4$ and to 4N has probability $1/4$. Going from 4K, 3K will occur with probability $1/3$ and 3N with probability $1/3$. From 3K, 2K will occur with probability $1/2$ and 2N with probability $1/2$. If the advantage is already lost, from 4N and 3N only 3N and 2N can result. The numbers inside the circles representing these situations are the probabilities for them. Starting with a probability of 1 for 5K, we go forward and compute all other probabilities.

In conclusion, the player who knows one answer to be wrong maintains her advantage in $3/4$ of the cases when there are only four answers left, in half of the cases when there are three answers left, and in $1/4$ of the cases if there are only two answers left.

The game tree has to be modified by including a random move before each move the clever player makes, provided she didn't lose the advantage before. The modified game trees can be found in the two cases where Ann, or Cindy have additional information on two sheets of the Excel Sheet [SeqQuizShow.xlsx](#). Let us discuss the results for one of the cases:

17.2.1 Cindy Knows one Answer to be False

In Figure 17.4, the expected payoffs are displayed, with the prize value on the horizontal axis. Beth's are largest values, followed by Ann and Cindy, but Cindy's expectations are now closer to Ann's, passing it for prize values larger than around 19. The sum of the expected values does not differ much from the case where nobody knows an answer to be wrong. For a prize value greater than 4 it is now 0.25 larger, and otherwise 0.25 smaller. So the gain in Cindy's expectation must be balanced somehow by losses to Ann and Beth. It turns out that both lose, but Beth more, and that both the gain for Cindy increases and the loss for Ann and Beth decrease for increasing prize value. The picture is in Figure 17.4.




Figure 17.4. Payoffs if Cindy knows one answer to be wrong

Exercises

Use the following facts for the exercises. In SEQUENTIAL QUIZ SHOW(9, 4), Ann's, Beth's, and Cindy's expectations are about 1.67, 3, 0.33. In the clever Ann variant they change to about 2.75, 2.25, 0.33. In the clever Beth variant they are about 1.67, 3, 0.33, and in the clever Cindy variant they are about 1.46, 2.63, 1.2.

- 1. Verify the backward induction analysis in Figure 17.1 by hand.

2. Assume the quiz master's assistant knows one of the answers to be wrong in advance. In SEQUENTIAL QUIZ SHOW(9, 4), to whom would the assistant reveal the wrong answer, and at what price?
3. With the same assumption as in the previous exercise, the assistant could also make money by promising not to reveal the information to anybody. How much money would he get, and from whom? This is an interesting option: Getting money for not cheating.

Project 30

SEQUENTIAL QUIZ SHOW, clever Ann For the clever Ann variant, create a graph like in Figures 17.2 and 17.4, showing prize value versus expected payoffs for Ann, Beth, and Cindy, in case there is a penalty of four units for a wrong answer. Describe thoroughly what happens.

Project 31

SEQUENTIAL QUIZ SHOW, clever Beth This is an exercise for students with good Excel skills: Modify the Excel sheet for a clever Beth, who knows one of the answers to be wrong from the beginning.

CHAPTER 18

Las Vegas 1962

In the 50s and 60s, Las Vegas was different from what it is now. In the 50s it still had a Wild West flavor: No carpets, no dress code, cowboy boots and hats worn in the casinos. In the 60s, when the Mafia took over many of the casinos, those on the strip became more elegant. Alcohol, illegal drugs, and prostitution were never far away, nor were cheating and violence.

Las Vegas then was a world far away from the academic world except for some mathematicians' interest in gambling. As discussed in Chapter 13, Probability theory started from questions about games of luck. Games of skill were discussed in Zermelo's paper on chess [Z1913], and in papers on simplified versions of poker by Borel, von Neumann, and Kuhn [K1950]. From then on, games, in most cases simplified, were taken as examples of simple games. Mathematicians, in particular the young Princeton graph theorists John Nash, Martin Shubik, Lloyd Shapley, and others played various games as Go, Hex, Kriegsspiel, extensively. Although the games are usually too complex for a complete analysis, their attractiveness for mathematicians may lie in the fact that there is no modeling necessary. The rules are firm, and the payoffs are given.

Most professional gamblers of that time were far away from the academic world, and they were not interested in what insight professors had to offer. The only mathematics they believed to be useful were the very simple comparison of odds and payoff ratio, which was standard and distinguished professional gamblers from the so-called "suckers". Many gamblers played in casinos and bet on sports events like horse racing. Their expertise was not in mathematics but in the sports involved—the key was to estimate the odds correctly.

Casino games are all games of chance—they all contain randomness and luck to some extent. You don't need a license to set up a chess tournament with a high prize for the winner, but you need a license for roulette, slot machines, blackjack, or poker. Still, the amount of luck involved differs. There is no skill in roulette but in poker, in the long run, skilled players will win against beginners. This is where game theory may be helpful.

During the 50s and 60s, mathematicians and mathematics started to influence casino play and gambling. The starting point was blackjack and the insight that both the odds as well as the best strategy depends on what's left in the deck of cards. Card counters try to estimate what's left by keeping track of the cards drawn. Jess Marcum, who was employed at RAND Corporation (see Chapter 21 for more on RAND) before becoming a professional gambler was one of the first to apply the technique. An early paper on card counting and the odds is [BCMM1957]. The optimal betting size was considered by J. L. Kelly Jr. [K1956]. Card counting and the Kelly criterion, were further developed and popularized by mathematics professor Edward O. Thorp.

Thorp wrote a paper where he improved the the results of [BCMM1957], using a computer for calculations. After presenting its contents in a talk with the promising title "Fortune's formula: A winning strategy for blackjack", a reporter became interested and interviewed Thorp. His article appeared in a local newspaper and Thorp got thousands of letters asking for the system, or offering to invest money. He accepted an offer from Emmanuel Kimmel, a bookmaker and gambler with mob connections. Together they went to Reno and

made thousands of dollars, using card counting and Kelly's criterion [P2005]. The winning spree in Reno got them a lot of attention from casino owners in Reno, but only in 1962, when Thorp published his best-seller book *Beat the Dealer* [T1962] which sold 700,000 copies, did gamblers and casinos become aware that mathematics could be useful for them.

Card counting is explained in Chapter 19. Here is the Kelly criterion:

Example Assume you are playing five rounds of a game where you win with probability 10% and lose with probability 90%. Assume that when winning you get your bet back, together with 100 times your bet (the **net odds** are 100). You have \$1000, and cannot get credit. How do you play? How much would you bet in each round?

The game is favorable to you. If you bet an amount of a in the first round, with 10% probability you will have $1000 + 100a$ after the first round, and with 90% probability you will have $1000 - a$. The expected value after one round is $0.1 \cdot (1000 + 100a) + 0.9 \cdot (1000 - a) = 1000 + 9.1a$. The larger the bet a , the larger the expected value of the money you have after each round. So you would always bet as much as you have, right?

Or rather not? When applying the strategy of always betting everything, when you bet the whole \$1000, with 90% probability all is gone after the first round and you can go home. Even if you win in the first round and bet all your money, the whole \$100,000, in the second round, it is very likely that you lose then. And so on. Playing this risky strategy, with probability $(0.1)^5 = 0.00001$, you win in all five rounds and take \$10,000,000,000,000 home, to become the richest human in the world, about 200 times as rich as Bill Gates. But almost surely, with probability 0.99999 you lose all your money at some point and go home without any money left. Still, the expected value for the game when playing the strategy is $0.00001 \cdot 10,000,000,000,000 + 0.99999 \cdot 0 = 100,000,000$. Not bad, but maybe you would prefer a somewhat higher probability to win a smaller amount of money. Wouldn't it be nicer to have \$1,000,000 after the five rounds, with probability 90%? However, this advantage is not reflected in the expected value, which is only \$900,000 for this option, much less than in the risky strategy.

The expected value approach may not be useful when large numbers are involved. I bet that most of us would prefer \$1,000,000 in hand to a 50-50 chance of winning \$3,000,000, though the second option has an expected value of \$1,500,000. The calculation is not flawed, but most people don't consider \$3,000,000 to be three times as valuable as \$1,000,000. An additional \$1000 is less valuable for a rich person than for a poor person. This problem can be overcome by introducing a **utility function**. The money itself is not considered to be the payoff, but the satisfaction $u(x)$ derived from possessing x dollars. Possible utility functions are the square root function $u(x) = \sqrt{x}$, where four times the money produces only twice the happiness, or $u(x) = \log(x + 1)$, where the logarithm (roughly) counts the number of digits of a number. The reason why we use $\log(x + 1)$ instead of $\log(x)$ is that $\log(x)$ is not defined for $x = 0$. Using $\log(x + 1)$ attaches a utility of $u(0) = \log(1) = 0$ to the value $x = 0$, and then slowly increasing utilities to increasing x . For instance, $u(9) = 1$, $u(99) = 2$, $u(999) = 3$, and so on.

Why choose the logarithm or the square root as the utility function? They have certain properties that utility is supposed to have: They are increasing—more money means more utility—but they increase less and less rapidly, which matches the observation that an increase of \$1 means less to somebody who has a lot than to somebody with little. The reason why the logarithm is often proposed as the utility function goes back to Daniel Bernoulli, who claimed that the increase of utility of a dollar would be inversely proportional to the amount of money already owned. An additional dollar is twice as valuable to somebody having \$1000 than to somebody having \$2000. This property is satisfied by the logarithm, as

those of you who have taken calculus should know. Though this claim is simple, and therefore attractive to mathematicians, it does not necessarily reflect reality.

The **Kelly criterion** tells how to bet to produce the highest expected utility using the logarithm as utility function in one round. Recall that in our example, if we bet a , with 10% probability you have the utility $\log(1000 + 100a + 1)$ after the first round, and with 90% probability you have the utility $\log(1000 - a + 1)$. So the expected utility is

$$0.1 \cdot \log(1001 + 100a) + 0.9 \cdot \log(1001 - a) = \log((1001 + 100a)^{0.1} \cdot (1001 - a)^{0.9}).$$

This value is largest if $(1001 + 100a)^{0.1} \cdot (1001 - a)^{0.9}$ is largest. Experiment with your calculator or with Excel to verify that $a = 91$ gives the largest value of about 1158 for it. The logarithm of 1158 value is about 3.06357. So according to the Kelly criterion you would bet \$91 in the first round. In subsequent rounds, the reasoning is similar, and you would always bet about 9.1% of your available money. (Note that Kelly used the function $\log(x)$ instead of $\log(x + 1)$, which yields a slightly different recommendation, still rounding to 9.1% for our parameters.)

The game can be simulated in the one choice sheet of the Excel sheet [Kelly.xlsx](#). In cell A36 you tell how much of your wealth you risk in each round, and see the probabilities for the outcomes to the right. If, for example, you risk 50% of your money in each round, you end, after five rounds, with

- \$318,781 with probability 7%
- \$3,156 with probability 33%
- \$31 with probability 59%

There are three more outcomes with low probabilities, namely \$328,440,640,656 with probability 0.001%, \$3,251,887,531 with probability 0.045%, and \$32,196,9061 with probability 0.81%. Neglecting them, you are most likely to lose almost everything, win about \$2156 (and to have \$3156), but there is also a real chance (7%) for a good win (\$317,781). If you risk 9.1% in each round, you will have after five rounds

- \$75.244 with probability 7%
- \$6.834 with probability 33%
- \$621 with probability 59%

Doesn't this look better? The two most likely outcomes now give a higher payoff.

The Kelly criterion depends on the choice of the utility function. If we select another one, as the square root function, we get a different conclusion in our example to bet about 58% of your money in each round.

The Kelly criterion can be used for all games where the player has an advantage. For games where there is a disadvantage, like roulette or slot machines, a rational player would better not play if the purpose is to win money. Poker is different, since you need to be a better player than the others to have an advantage, and most people think that they are above average.

Back to Thorp and the 60s. The other technique popularized in Thorp's book, card counting in blackjack, became more and more an issue in the casinos. Eventually casinos took countermeasures by changing the rules: More decks of cards were used in the shoe. The shoe was shuffled more frequently. Limits on the maximum bet (or better, the ratio of maximum and minimum bet) were introduced. Although using your brain cannot reasonably be forbidden, card counters were banned from casinos, and lists of card counters were maintained. Card counters had to reveal themselves by sudden changes between low and very high bets, so they started to work in groups, with someone betting always high, but going to a table when signaled by a

low bet player that the odds are favorable. Ken Uston was successful with this method, and an MIT student group worked in this way during the 80s and 90s. The movie “21” is based on their story.

Much has changed since the 60s. Las Vegas is now family-friendly and non-smoking. There may still be drugs or organized crime, but they are not visible on the surface. Online poker and other online games have become popular. The casino’s and the professional gambler’s opinion towards mathematics has changed. Gamblers know and appreciate what game theory can offer to blackjack and to poker and other games. Though poker has not been mathematically solved, some mathematicians create competitive poker programs based on game theory. One of the best players, Chris (“Jesus”) Ferguson, has a Ph.D. degree in computer science and publishes papers on game theory.

Some authors claim that card counters still have a slight advantage. And why not? It would be rational for the casinos to keep alive the rumor that it is possible to win. More people would come and play, most of them would not be able to play optimally and would lose.

Exercises

1. Use the one choice sheet in the Excel sheet [Kelly.xlsx](#) to produce a graph displaying the expected utility with the logarithmic utility function (on the y -axis) against the percentage of money bet (on the x -axis). Use at least five percentages. Do another graph displaying the expected value of money won depending on the percentage of money bet.
2. Use the one choice sheet in the Excel sheet [Kelly.xlsx](#) to produce a graph displaying the expected utility with the square root utility function (on the y -axis) against the percentage of money bet (on the x -axis). Use at least five percentages. Do another graph displaying the expected value of money won depending on the percentage of money bet.
3. What changes if the net odds (the amount of dollars you win for every dollar of your bet in case you win) of the game are only 10? Use the one choice sheet in the Excel sheet [Kelly.xlsx](#) to produce a graph, displaying the utility with the logarithmic utility function against the the percentage of money bet. What is the best betting percentage then?
4. If you know that you will play only five rounds, should you adjust the percentage of money you bet depending on the round? Use the many choices sheet in the Excel sheet [Kelly.xlsx](#) and report whether you can improve the expected logarithmic utility achieved with a constant betting percentage of 9.1%.
5. Assume you play the five rounds betting game and know you need to have \$10000 at the end, otherwise you are dead (since you owe that much to the mob). Would you play differently from the 9.1% strategy? How? Could you play so that you have at least a 30% probability of surviving? How large could you make this probability? Use the many choices sheet in the Excel sheet [Kelly.xlsx](#).
6. In the five round betting model in the Excel sheet [Kelly.xlsx](#), there are $2^5 = 32$ outcomes. But if the percentage of money used is the same in each round, as in the one choice sheet, there are only six possible money amounts. For instance, if you win, win, lose, win, and lose you have at the end the same amount of money as if you lose, lose, win, win, and win. Explain why this is so, and why it is not so when the player decides in each round and in each position how much to bet, as in the many choices sheet.

CHAPTER 19

Example: Mini Blackjack and Card Counting

Prerequisites: Chapters 8, 12, and 16.

Do you gamble at a casino? Would you? Why not? If your answer is that casino games are designed so that the Casino has better odds, you are only partially right. Blackjack is one of the few casino games where playing optimally may actually win you money in the long run. In the last chapter we mentioned some attempts to analyze the game. For those of you eager to go and bankrupt a casino, you will be disappointed to hear that we will not discuss blackjack in detail. There are two reasons. First, casino blackjack is too complicated to treat in an introductory book. Second, and more important, it has a huge extensive form. A similar game called MINI BLACKJACK has most of the essential ingredients of casino blackjack, and it is simple enough for us to analyze. We will see three aspects of casino blackjack, namely counting cards, playing against the house, and reshuffling, in versions of the game.

19.1 The Basic Game

The basic game depends on two parameters a and b :

MINI BLACKJACK(a, b) This two-person zero-sum game is played with a shuffled deck of cards containing a known number a of 1s and a known number b of 2s. Ann and Beth, get a card from the top of the deck in front of them face up. A player makes a move by choosing to draw no more cards or to receive another card from the top of the deck. If the player draws another card, it is shown face up. Beginning with Ann, the two players alternate moves unless a player has already chosen to stop drawing cards, in which case the player cannot move anymore. The game ends when both players have stopped. The goal of the game is to have the sum of the player's own cards be 3 or as close to 3 as possible, but not greater than 3. Thus a sum of 3 is best, better than a sum of 2, which is better than a sum of 1, which is better than all other sums (4, 5, or higher). The payoff is 1 for the winner, -1 for the loser, and 0 for a draw.

Student Activity Play ten rounds of MINI BLACKJACK(16, 16) against your neighbor to learn how to play. If you don't have cards, you can play the game with two players in the applet [MiniBlackJack](#).

Because the game is sequential game with perfect information, it can be displayed using an extensive form and can be analyzed using backward induction. Let us start with the extensive form, for which we will use a game tree instead of a game digraph. The positions where Ann has to move are uniquely determined by the cards Ann and Beth have so far and by whom has stopped. This is encoded by writing the values of Ann's

cards, then a hyphen, and then the values of Beth’s cards. If Ann has stopped, then the string of her cards is followed by a period, and the same for Beth. For example, 11.-12 means that Ann has two cards of value 1 each and has stopped, and that Beth has one card of value 1 and one of value 2. We use the same encoding for positions where Beth is to move. We indicate whose turn it is to move by the color of the vertex and by labels.

For each choice of a and b we have a different game tree. Fortunately its essential structure remains the same. Only the probabilities for the random moves depend on a and b . For instance, position 1-1 is more likely in the version starting with $a = 8$ 1s and $b = 4$ 2s than in the version starting with $a = 4$ 1s and $b = 8$ 2s (to be precise, it occurs with probability $\frac{a}{a+b} \frac{a-1}{a+b-1}$).

This structure without the probabilities for the random moves is shown in Figure 19.2.

We have omitted from the game tree some obviously bad options. For instance, when a player has only one card, a 1, it doesn’t make sense to stop (see Ann’s (red) 1-1 or 1-2 positions). Also, if a player is behind, stopping doesn’t make sense either (see Ann’s (red) 11-12 position). And if a player has 3 points, she should stop (see Ann’s (red) 12-11 position).

To analyze the game tree, we open the Excel sheet `MiniBlackJack.xlsx` at the sheet “Game”. It contains the game tree, and it updates the probabilities for the random moves when the user types the numbers a and b into the two grey fields. The backward induction is done automatically and the blue and green cells under the leftmost vertex indicate the expected payoffs for Ann and Beth in the corresponding version MINI BLACKJACK(a, b). When trying it out, you will notice that the expected payoff for Ann is usually negative.

I hope the reader agrees with me that the process of backward induction, though tedious and time-consuming, is in principle easy, easy enough for an Excel sheet to solve the game. But Excel can do more: It can compute all expected payoffs in the solutions for Ann and Beth for all possible combinations of a and b .

The data are shown in the table at the bottom of the sheet “Game” and in Figure 19.1.

		2s																
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1s	0	####	####	####	####	####	####	0	0	0	0	0	0	0	0	0	0	0
	1	####	####	####	####	####	0	0	0	0	0	0	0	0	0	0	0	0
	2	####	####	####	####	-0.07	-0.05	-0.04	-0.03	-0.02	-0.02	-0.02	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01
	3	####	####	####	-0.05	-0.09	-0.05	-0.05	-0.04	-0.04	-0.03	-0.03	-0.02	-0.02	-0.02	-0.02	-0.02	-0.01
	4	####	####	0	-0.03	-0.06	-0.09	-0.05	-0.05	-0.04	-0.04	-0.04	-0.03	-0.03	-0.03	-0.03	-0.02	-0.02
	5	####	0	0	0	-0.04	-0.06	-0.07	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.03	-0.03	-0.03	-0.03
	6	0	0	0	0	-0	-0.05	-0.07	-0.07	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.03	-0.03
	7	0	0	-0.03	-0	-0	-0.02	-0.05	-0.07	-0.06	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.03
	8	0	0	-0.02	-0.03	-0	-0.01	-0.03	-0.06	-0.07	-0.06	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04
	9	0	0	-0.02	-0.03	-0.02	-0	-0.01	-0.04	-0.06	-0.06	-0.05	-0.04	-0.04	-0.04	-0.04	-0.04	-0.04
	10	0	0	-0.02	-0.02	-0.03	-0.02	-0.01	-0.02	-0.04	-0.07	-0.06	-0.05	-0.04	-0.04	-0.04	-0.04	-0.04
	11	0	0	-0.01	-0.02	-0.03	-0.03	-0.02	-0.01	-0.03	-0.05	-0.07	-0.06	-0.05	-0.04	-0.04	-0.04	-0.04
	12	0	0	-0.01	-0.02	-0.03	-0.03	-0.03	-0.02	-0.01	-0.03	-0.05	-0.06	-0.06	-0.05	-0.04	-0.04	-0.04
	13	0	0	-0.01	-0.02	-0.02	-0.03	-0.03	-0.03	-0.02	-0.02	-0.04	-0.05	-0.06	-0.05	-0.05	-0.04	-0.04
	14	0	0	-0.01	-0.02	-0.02	-0.03	-0.03	-0.03	-0.02	-0.02	-0.02	-0.04	-0.06	-0.06	-0.05	-0.04	-0.04
	15	0	0	-0.01	-0.01	-0.02	-0.02	-0.03	-0.03	-0.03	-0.02	-0.02	-0.02	-0.03	-0.04	-0.06	-0.06	-0.05
	16	0	0	-0.01	-0.01	-0.02	-0.02	-0.03	-0.03	-0.03	-0.03	-0.03	-0.03	-0.02	-0.03	-0.05	-0.06	-0.06

Figure 19.1. Expected payoffs for Ann for one play of MINI BLACKJACK(a, b), where a is the number of 1s and b the number of 2s in the deck

For instance, if you play MINI BLACKJACK(7, 5) with seven 1s and five 2s, the expected payoff for Ann at the beginning is -0.02. The more purple the cells, the greater the advantage for Beth. The ##### cells indicate that this instance of the game cannot be played. For example, it makes no sense to play the game with only five cards in the deck, since then it may get empty during the play.

The probabilities of the random moves, and the expected payoffs change when a and b change, and the strategies of the players might also change. Here is an example. Look at the situation (a blue cell, labeled by




Figure 19.2. The game tree

11-2. in the Excel sheet) where Ann has two cards of value 1 each, and Beth has only one card, a 2, and has stopped. The question is whether Ann should take another card. Stopping means a draw with payoff of 0. If $a = 5$ and $b = 2$, then since Ann got two of these 1s and Beth got one of the 2s, there are three 1s and three 2s left. If Ann asks for another card, then Ann wins with probability $\frac{3}{6}$ —if the next card is a 1—but loses otherwise, with probability $\frac{3}{6}$. So asking for another card has the expected payoff of $\frac{3}{6} \cdot 1 + \frac{3}{6} \cdot (-1) = 0.5$ for Ann, and is therefore preferable to stopping. If $a = 5$ and $b = 3$, the expected payoff if drawing another card is still 0.2, but it is 0 for $a = 5$ and $b = 4$ and -0.14 for $a = 5$ and $b = 5$. For a less obvious example, look at the situation (a blue cell, labeled 2-2 on the Excel sheet) where both players have a 2 and Ann has to decide whether to get another card. If $a = 5$ and $b = 5$, stopping has an expected payoff of -0.25 , but getting another card an expected payoff of -0.11 . Therefore Ann would get another card. However, if $a = 5$ and $b = 6$, stopping has an expected payoff of -0.11 , and getting another card an expected payoff of -0.17 , so Ann would stop. To be able to play differently, the players must know a and b .

All expectations for Ann are negative or zero. Beth, moving second, has the advantage.

In reality, we are not interested in playing these variants. They are part of the game. **MINI BLACKJACK** without parameters. It starts with a play of MINI BLACKJACK(16, 16). After the first play, a 1s and b 2s remain in the deck, then a round of MINI BLACKJACK(a , b) follows, and so on. There is no reshuffling after each play. Theoretically we also need a termination criterion, say a fixed number of rounds.

19.2 Playing against the House

In casino blackjack, players are playing against the house, represented by the dealer. The dealer may not change strategy depending on the distribution of the cards still in the deck. The dealer decides whether to get a new card by fixed rules. In casino blackjack, the disadvantage of having to play with less flexibility is offset by the dealer's moving last. Since the dealer has to play mechanically anyway, this is not so much an advantage for the dealer as it is a disadvantage for the other players, as they have not seen the dealer's draw when they have to make their move.

To model the dealer's way of playing in MINI BLACKJACK, we say that the dealer has to play as he or she would play optimally in MINI BLACKJACK(16, 16). We make Beth in the second position (with her slight advantage) the dealer.

We start with 16 1s and 16 2s, and continue several plays without reshuffling. Dealer Beth has to take another card if she is behind (11-1, 12-1, 2.-1, 12.-1, 111-11, 12.-11, 12-2, 12-2). Otherwise she gets a card if both have one card, a 2 each, and Ann has stopped (2.-2). The dealer stops in the remaining cases of a draw (11.-11, 2.-11, 11-2).

Ann can adjust her play to the situation. To do that she needs to count cards, keeping track of the cards drawn so far. A card counting player knows the number of 1s and 2s still in the deck. A card counter plays differently for different a and b , and also, as we will see later, by betting differently.

Student Activity Play twenty rounds against the computer dealer Beth in the applet [MiniBlackJack1](#).

The Dealer sheet on the Excel sheet models the modified game. The dealer doesn't have any options for her "moves" anymore, so the game could also be modeled as a 1-person game. Different numbers a and b result in different expected payoffs for Ann, and Figure 19.3 displays Ann's payoffs (assuming best play for Ann) for the distributions of 1s and 2s.

Let us look at $a = 6$ and $b = 2$. There are four distributions: A 1 for both players, a 1 for Ann and a 2 for Dealer Beth, a 2 for Ann and a 1 for Dealer Beth, and a 2 for both players. They have probabilities $\frac{6}{8} \cdot \frac{5}{7} = \frac{15}{28}$, $\frac{6}{8} \cdot \frac{2}{7} = \frac{6}{28}$, $\frac{2}{8} \cdot \frac{6}{7} = \frac{6}{28}$, and $\frac{2}{8} \cdot \frac{1}{7} = \frac{1}{28}$. Look at sheet "Dealer" in the Excel sheet [MiniBlackJack.xlsx](#), and type 6 into cell E11 and 2 into cell E12. Backward analysis shows that Ann's expected payoff is 0 in the first case, $\frac{1}{2}$ in the second case, and 0 in the third and fourth cases. That means that Ann's expected payoff in MINI BLACKJACK(6, 2) is $\frac{6}{28} \cdot \frac{1}{2} = \frac{3}{28} \approx 0.107$.

		2s																	
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
1s	0	####	####	####	####	####	####	1	1	1	1	1	1	1	1	1	1	1	
	1	####	####	####	####		0.333	0.429	0.5	0.556	0.6	0.636	0.667	0.692	0.714	0.733	0.75	0.765	
	2	####	####	####	####		-0.07	0.048	0.143	0.222	0.289	0.345	0.394	0.436	0.473	0.505	0.533	0.559	0.582
	3	####	####	####		-0.05	-0.09	-0.05	0.012	0.075	0.133	0.186	0.234	0.277	0.316	0.352	0.384	0.413	0.44
	4	####	####	0.067	-0.03	-0.06	-0.09	-0.05	-0	0.042	0.087	0.129	0.168	0.205	0.24	0.272	0.302	0.329	
	5	####	0.167	0.095	0.018	-0.04	-0.06	-0.07	-0.04	-0.01	0.024	0.059	0.093	0.126	0.158	0.188	0.216	0.243	
	6	0	0.143	0.107	0.048	-0	-0.05	-0.07	-0.07	-0.04	-0.02	0.012	0.041	0.069	0.097	0.124	0.151	0.176	
	7	0	0.125	0.111	0.067	0.021	-0.02	-0.05	-0.07	-0.06	-0.04	-0.02	0.004	0.028	0.052	0.076	0.1	0.123	
	8	0	0.111	0.111	0.079	0.04	0.004	-0.03	-0.06	-0.07	-0.06	-0.04	-0.02	-0	0.019	0.04	0.061	0.082	
	9	0	0.1	0.109	0.086	0.055	0.022	-0.01	-0.04	-0.06	-0.06	-0.05	-0.04	-0.02	-0.01	0.012	0.031	0.049	
	10	0	0.091	0.106	0.091	0.065	0.036	0.008	-0.02	-0.04	-0.07	-0.06	-0.05	-0.04	-0.02	-0.01	0.007	0.023	
	11	0	0.083	0.103	0.093	0.073	0.048	0.022	-0	-0.03	-0.05	-0.07	-0.06	-0.05	-0.04	-0.02	-0.01	0.003	
	12	0	0.077	0.099	0.095	0.078	0.057	0.034	0.011	-0.01	-0.03	-0.05	-0.06	-0.06	-0.05	-0.04	-0.03	-0.01	
	13	0	0.071	0.095	0.095	0.082	0.064	0.043	0.022	0.002	-0.02	-0.04	-0.05	-0.06	-0.05	-0.05	-0.04	-0.03	
	14	0	0.067	0.092	0.094	0.085	0.069	0.051	0.032	0.013	-0.01	-0.02	-0.04	-0.06	-0.06	-0.05	-0.04	-0.04	
	15	0	0.063	0.088	0.093	0.086	0.073	0.057	0.04	0.022	0.004	-0.01	-0.03	-0.04	-0.06	-0.06	-0.05	-0.04	
	16	0	0.059	0.085	0.092	0.088	0.077	0.063	0.047	0.03	0.014	-0	-0.02	-0.03	-0.05	-0.06	-0.06	-0.05	

Figure 19.3. Expected payoff for Ann when playing one game of MINI BLACKJACK(*a*, *b*) (with *a* 1s and *b* 2s) against dealer Beth

Red cells indicate an advantage for the dealer, Beth, and green cells an advantage for Ann. Some card distributions are very favorable for Ann, like those with many 2s (as $b \geq 13$) and few 1s (as $a \leq 2$) with expected payoff larger than 0.5 for Ann. Other card distributions in the deck, like $a = 8$ and $b = 8$ with expected payoff of -0.07 for Ann, are slightly favorable for the dealer. The more the distribution in the deck deviates from an equal distribution of 1s and 2s, the better for Ann. Probably this is because Beth’s fixed strategy is designed to work for $a = 16$ and $b = 16$. In addition, when there are very few 2s it is always best to take another card, but the dealer’s mechanical strategy would not always call for this.

From Figure 19.3, you may think that since there are more cells favorable to Ann, she has the advantage. The average of all entries in the table in Figure 19.3 is 0.116, so doesn’t Ann have an advantage? The answer is no, and the reason is that the cells are not equally likely to occur. It is more likely to have a round of the game with eight 1s and eight 2s than one with sixteen 1s and no 2s left. The more the card distribution deviates from equality between 1s and 2s, the rarer it is. The rare situations are the ones where Ann has an advantage. This has to be taken into account.

19.2.1 How Likely are the Distributions?

To get a better idea of Ann’s average chances, we will estimate the likelihood of the subgames MINI BLACKJACK(*a*, *b*) during a sequence of several rounds, starting with a deck of 16 1s and 16 2s. The numbers are difficult to compute,so we instead look at how often different card distributions in the deck occur each time a card is dealt (not just at the beginning of the rounds).

We use a probability digraph. Let *a*-*b* denote the distribution of *a* 1s and *b* 2s. Since we start with 16-16, this occurs with probability 1. When removing one card, we get a 1 or a 2 with the same probability $16/32 = 1/2$, so 15-16 and 16-15 have the same probability of $1/2$. When we remove the next card, some distributions become more likely than others. Since 16-15 would result in 16-14 and 15-15 with about equal probability (of $15/31$ and $16/31$), and since 15-16 would result in 15-15 and 14-16 with about equal probability, 14-16 and 16-14 occur with roughly probability $1/4$, whereas 15-15 occurs with probability of about $1/2$. The beginning of the probability digraph is in Figure 19.4 The probabilities of the arcs are at the arcs, and the probabilities of the situations at the vertices. The simple but tedious calculations are left to the Excel sheet [CardCounting.xlsx](#). The numbers are in the table in Figure 19.5. The start of the probability digraph is in the lower right corner of the table, and the movement is up and to the left.




Figure 19.4. Start of the probability digraph

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	1.0000	0.5000	0.2419	0.1129	0.0506	0.0217	0.0088	0.0034	0.0012	0.0004	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1	0.5000	0.5161	0.3871	0.2492	0.1446	0.0771	0.0381	0.0174	0.0073	0.0028	0.0010	0.0003	0.0001	0.0000	0.0000	0.0000	0.0000
2	0.2419	0.3871	0.4004	0.3337	0.2410	0.1557	0.0914	0.0489	0.0239	0.0106	0.0043	0.0015	0.0005	0.0001	0.0000	0.0000	0.0000
3	0.1129	0.2492	0.3337	0.3461	0.3028	0.2326	0.1599	0.0993	0.0559	0.0284	0.0129	0.0052	0.0018	0.0005	0.0001	0.0000	0.0000
4	0.0506	0.1446	0.2410	0.3028	0.3149	0.2834	0.2259	0.1614	0.1037	0.0599	0.0309	0.0141	0.0055	0.0018	0.0005	0.0001	0.0000
5	0.0217	0.0771	0.1557	0.2326	0.2834	0.2957	0.2711	0.2213	0.1618	0.1060	0.0618	0.0317	0.0141	0.0052	0.0015	0.0003	0.0000
6	0.0088	0.0381	0.0914	0.1599	0.2259	0.2711	0.2840	0.2637	0.2186	0.1619	0.1067	0.0618	0.0309	0.0129	0.0043	0.0010	0.0001
7	0.0034	0.0174	0.0489	0.0993	0.1614	0.2213	0.2637	0.2776	0.2603	0.2177	0.1619	0.1060	0.0599	0.0284	0.0106	0.0028	0.0004
8	0.0012	0.0073	0.0239	0.0559	0.1037	0.1618	0.2186	0.2603	0.2756	0.2603	0.2186	0.1618	0.1037	0.0559	0.0239	0.0073	0.0012
9	0.0004	0.0028	0.0106	0.0284	0.0599	0.1060	0.1619	0.2177	0.2603	0.2776	0.2637	0.2213	0.1614	0.0993	0.0489	0.0174	0.0034
10	0.0001	0.0010	0.0043	0.0129	0.0309	0.0618	0.1067	0.1619	0.2186	0.2637	0.2840	0.2711	0.2259	0.1599	0.0914	0.0381	0.0088
11	0.0000	0.0003	0.0015	0.0052	0.0141	0.0317	0.0618	0.1060	0.1618	0.2213	0.2711	0.2957	0.2834	0.2326	0.1557	0.0771	0.0217
12	0.0000	0.0001	0.0005	0.0018	0.0055	0.0141	0.0309	0.0599	0.1037	0.1614	0.2259	0.2834	0.3149	0.3028	0.2410	0.1446	0.0506
13	0.0000	0.0000	0.0001	0.0005	0.0018	0.0052	0.0129	0.0284	0.0559	0.0993	0.1599	0.2326	0.3028	0.3461	0.3337	0.2492	0.1129
14	0.0000	0.0000	0.0000	0.0001	0.0005	0.0015	0.0043	0.0106	0.0239	0.0489	0.0914	0.1557	0.2410	0.3337	0.4004	0.3871	0.2419
15	0.0000	0.0000	0.0000	0.0000	0.0001	0.0003	0.0010	0.0028	0.0073	0.0174	0.0381	0.0771	0.1446	0.2492	0.3871	0.5161	0.5000
16	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0004	0.0012	0.0034	0.0088	0.0217	0.0506	0.1129	0.2419	0.5000	1.0000

Figure 19.5. Theoretical relative frequencies of the distributions of 1s and 2s in the deck when starting with 16 1s and 16 2s and drawing cards repeatedly

Let’s clarify our game MINI BLACKJACK (against the dealer). We start with a deck of sixteen 1s and sixteen 2s, and play MINI BLACKJACK(16, 16). After this, we play MINI BLACKJACK(a , b) with the remaining deck of a 1s and b 2s. Each play played is called a round. We proceed until, at the end of some play or round, the cards are collected and reshuffled, and we play MINI BLACKJACK(16, 16) again, continue with more rounds, reshuffle again, and so on.

Since reshuffling is only done after a round is finished, the numbers in Figure 19.5 do not reflect how often the different versions of MINI BLACKJACK(a , b) occur as a play. Since at every round 2 to 6 cards are dealt, no round would start with distribution 16-15, for instance. The entry 0.3337 in the cell (14,13) in the table in Figure 19.5 is not the probability that a round starts with 14 1s and 13 2s, but it is the probability that a game starting with a full deck of 32 cards has 14 1s and 13 2s at some point.

What we want are the probabilities that MINI BLACKJACK(a , b) occurs as a round. For simplicity, we will use the relative frequencies in Figure 19.5.

2s		Note that all table values in this table are given in 1/1000s															
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1s	0						0.3273	0.1259	0.0453	0.0151	0.0046	0.0013	0.0003	6E-05	9E-06	1E-06	6E-08
	1					0.9521	0.6042	0.3223	0.1511	0.0631	0.0234	0.0076	0.0021	0.0005	9E-05	1E-05	8E-07
	2				-0.595	0.2747	0.4834	0.4028	0.2561	0.1361	0.0621	0.0244	0.0081	0.0022	0.0005	7E-05	5E-06
	3			-0.641	-0.961	-0.461	0.0705	0.2758	0.2758	0.1958	0.112	0.0533	0.0211	0.0068	0.0017	0.0003	3E-05
	4		0.5951	-0.32	-0.666	-0.916	-0.398	-0.018	0.163	0.1925	0.1476	0.0877	0.0419	0.016	0.0047	0.0009	1E-04
	5	0.4761	0.5493	0.1538	-0.417	-0.695	-0.739	-0.362	-0.065	0.0941	0.135	0.1095	0.0657	0.0303	0.0105	0.0025	0.0003
	6	0	0.2014	0.3625	0.282	-0.04	-0.478	-0.717	-0.638	-0.34	-0.093	0.0488	0.0936	0.0794	0.0465	0.0196	0.0055
	7	0	0.0806	0.2014	0.2452	0.1268	-0.155	-0.524	-0.737	-0.575	-0.326	-0.112	0.0176	0.063	0.0552	0.0301	0.0105
	8	0	0.0302	0.0985	0.163	0.1552	0.0233	-0.235	-0.562	-0.707	-0.536	-0.318	-0.125	-0.004	0.0399	0.0355	0.0166
	9	0	0.0105	0.043	0.0908	0.1211	0.0863	-0.052	-0.295	-0.598	-0.665	-0.514	-0.314	-0.134	-0.02	0.0226	0.0197
	10	0	0.0033	0.0167	0.0435	0.0744	0.0831	0.0326	-0.109	-0.344	-0.635	-0.643	-0.505	-0.315	-0.14	-0.029	0.01
	11	0	0.001	0.0057	0.0179	0.0377	0.056	0.0507	-0.01	-0.154	-0.388	-0.665	-0.638	-0.508	-0.319	-0.141	-0.032
	12	0	0.0002	0.0017	0.0063	0.0159	0.0294	0.0385	0.0241	-0.042	-0.19	-0.429	-0.666	-0.652	-0.525	-0.326	-0.135
	13	0	5E-05	0.0004	0.0018	0.0055	0.0122	0.0206	0.0232	0.0033	-0.067	-0.218	-0.47	-0.684	-0.691	-0.563	-0.333
	14	0	8E-06	8E-05	0.0004	0.0015	0.0039	0.008	0.0125	0.0111	-0.011	-0.081	-0.237	-0.513	-0.728	-0.775	-0.637
	15	0	1E-06	1E-05	7E-05	0.0003	0.0008	0.0021	0.0042	0.006	0.0029	-0.018	-0.083	-0.24	-0.524	-0.818	-0.97
	16	0	6E-08	8E-07	5E-06	3E-05	1E-04	0.0003	0.0007	0.0014	0.0017	-8E-04	-0.015	-0.063	-0.201	-0.497	-1.028

Figure 19.6. Products of expected values and the probabilities of encountering the distributions of 1s and 2s at the beginning of a round, divided by 27, shown as 1/1000s

The sum of the values in Figure 19.5 is 32. If we rescale them, dividing numbers by 32, we get the probability that a round starts with that many 1s and 2s. Since we reshuffle at the end of one round if there are five or fewer cards in the deck, and therefore disregard the upper right corner of Figure 19.5, we divide by 27. Multiplying the rescaled numbers by the expected payoffs in Figure 19.3 for Ann provided we start the round with a deck with that many 1s and 2s, we get the values in Figure 19.6. Their sum is Ann’s expected payoff in a random round. It is -0.027 . Ann loses about 3 cents per round when each round is played for \$1. It is better than roulette, but it is still a loss in the long run.

19.2.2 Betting High and Low

We have seen that Ann can count cards and adjust her play, and dealer Beth has a fixed strategy. Still Ann has a slight disadvantage. But she can do something. There is another asymmetry between Ann and dealer Beth that we can take into account. In every round, Ann decides in advance how much to bet and Beth must match her bet.

This leads to a simple winning strategy for Ann. Ann bets nothing unless a round starts with a distribution favorable for her. However, this is not allowed: if you sit at the table, you must play. There is a minimum bet and a maximum bet. What Ann could do is to bet low if her expectation is negative, and high if it is positive. Ann counts cards for two reasons: to adjust her playing strategy, and to know how much she should bet.

For example, assume that four rounds are played and that the four subgames have expected payoffs of -0.07 , -0.03 , 0.1 , and -0.06 . Although the positive expectation is larger than the negative ones, their sum is negative. But if Ann bets \$1 in each of the negative rounds, but \$2 in the positive round, her expected value is $-0.07 \cdot 1 - 0.03 \cdot 1 + 0.1 \cdot 2 - 0.06 \cdot 1 = 0.04$. By betting higher in the round of positive expectations Ann gets a positive expected payoff. Betting low in negative expectation rounds and high in positive expectation rounds is only limited by the casino’s rules about minimum and maximum bets, and by the player’s desire not be revealed as a card counter, since she might be banned.

Ann will play with maximum bet n when she has an advantage, and play with minimum bet 1 otherwise. Suppose $n = 20$. The table in Figure 19.7 shows the products of Ann’s bets and her expected payoffs from the table in Figure 19.6. The negative values remain, and the positive values are multiplied by $n = 20$. The sum of the values is Ann’s expectation in a random round when betting correctly. It is about $197/1000 = 0.197$. That means that Ann will win on average 20 cents per round.

Student Activity Play against the computer dealer Beth with decisions on bets in the applet [MiniBlackJack2](#).

n=	20	2s	Note that all table values in this table are given in 1/1000s															
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1s	0							6.5459	2.5177	0.9064	0.3021	0.0919	0.0251	0.006	0.0012	0.0002	2E-05	1E-06
	1						19.043	12.085	6.4452	3.0212	1.261	0.4681	0.1529	0.043	0.0101	0.0018	0.0002	2E-05
	2					-0.595	5.4931	9.6678	8.0565	5.1229	2.7227	1.2419	0.4872	0.1622	0.0444	0.0095	0.0014	0.0001
	3				-0.641	-0.961	-0.461	1.4099	5.517	5.517	3.9168	2.2402	1.0664	0.4224	0.136	0.0338	0.0058	0.0005
	4			11.902	-0.32	-0.666	-0.916	-0.398	-0.018	3.26	3.8499	2.9512	1.7539	0.8388	0.3202	0.0933	0.0187	0.002
	5		9.5213	10.986	3.0761	-0.417	-0.695	-0.739	-0.362	-0.065	1.8825	2.6995	2.1908	1.3138	0.6065	0.2099	0.0496	0.0061
	6	0	4.0283	7.2509	5.6396	-0.04	-0.478	-0.717	-0.638	-0.34	-0.093	0.977	1.8726	1.5876	0.9301	0.3923	0.1108	0.0162
	7	0	1.6113	4.0283	4.904	2.5356	-0.155	-0.524	-0.737	-0.575	-0.326	-0.112	0.3528	1.26	1.1032	0.6027	0.2103	0.0372
	8	0	0.6042	1.9703	3.26	3.1048	0.4657	-0.235	-0.562	-0.707	-0.536	-0.318	-0.125	-0.004	0.7983	0.7108	0.3315	0.074
	9	0	0.2102	0.8598	1.8151	2.4217	1.7256	-0.052	-0.295	-0.598	-0.665	-0.514	-0.314	-0.134	-0.02	0.451	0.3949	0.1234
	10	0	0.0669	0.3344	0.8693	1.487	1.6624	0.6513	-0.109	-0.344	-0.635	-0.643	-0.505	-0.315	-0.14	-0.029	0.2006	0.1521
	11	0	0.0191	0.1146	0.359	0.755	1.1196	1.014	-0.01	-0.154	-0.388	-0.665	-0.638	-0.508	-0.319	-0.141	-0.032	0.0462
	12	0	0.0048	0.0339	0.1261	0.3185	0.5888	0.7698	0.4811	-0.042	-0.19	-0.429	-0.666	-0.652	-0.525	-0.326	-0.135	-0.025
	13	0	0.001	0.0084	0.0366	0.1093	0.2445	0.412	0.465	0.0659	-0.067	-0.218	-0.47	-0.684	-0.691	-0.563	-0.333	-0.108
	14	0	0.0002	0.0016	0.0083	0.029	0.0773	0.1605	0.25	0.2225	-0.011	-0.081	-0.237	-0.513	-0.728	-0.775	-0.637	-0.32
	15	0	2E-05	0.0002	0.0013	0.0054	0.0168	0.0422	0.0838	0.1199	0.0573	-0.018	-0.083	-0.24	-0.524	-0.818	-0.97	-0.806
	16	0	1E-06	2E-05	0.0001	0.0005	0.0019	0.0058	0.0141	0.0275	0.0346	-8E-04	-0.015	-0.063	-0.201	-0.497	-1.028	-1.833

Figure 19.7. Products of expected payoffs for Ann when betting correctly and frequencies. Values are in 1/1000s

19.2.3 Reshuffling

In casino blackjack, the deck is reshuffled frequently. Often six to eight decks of cards are used in a shoe, which is reshuffled as soon as about two thirds or three fourths of the cards have been dealt. Some casinos shuffle more often, some after each round.

For our smaller game, assume the cards are reshuffled at the end of a round if only a fixed number $k \geq 5$ of cards is left in the deck. Suppose $k = 15$. Then none of the values in the upper triangle is relevant, since no round will be started with those combinations of 1s and 2s. In the table in Figure 19.8, those values are eliminated.

n=	20	2s	Note that all table values in this table are given in 1/1000s																	
		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16		
1s	0																	1E-06		
	1																0.0002	2E-05		
	2																0.0095	0.0014		
	3														0.136	0.0338	0.0058	0.0005		
	4													0.8388	0.3202	0.0933	0.0187	0.002		
	5												2.1908	1.3138	0.6065	0.2099	0.0496	0.0061		
	6											0.977	1.8726	1.5876	0.9301	0.3923	0.1108	0.0162		
	7										-0.326	-0.112	0.3528	1.26	1.1032	0.6027	0.2103	0.0372		
	8										-0.707	-0.536	-0.318	-0.125	-0.004	0.7983	0.7108	0.3315	0.074	
	9										-0.295	-0.598	-0.665	-0.514	-0.314	-0.134	-0.02	0.451	0.3949	0.1234
	10							0.6513	-0.109	-0.344	-0.635	-0.643	-0.505	-0.315	-0.14	-0.029	0.2006	0.1521		
	11						1.1196	1.014	-0.01	-0.154	-0.388	-0.665	-0.638	-0.508	-0.319	-0.141	-0.032	0.0462		
	12					0.3185	0.5888	0.7698	0.4811	-0.042	-0.19	-0.429	-0.666	-0.652	-0.525	-0.326	-0.135	-0.025		
	13				0.0366	0.1093	0.2445	0.412	0.465	0.0659	-0.067	-0.218	-0.47	-0.684	-0.691	-0.563	-0.333	-0.108		
	14			0.0016	0.0083	0.029	0.0773	0.1605	0.25	0.2225	-0.011	-0.081	-0.237	-0.513	-0.728	-0.775	-0.637	-0.32		
	15		2E-05	0.0002	0.0013	0.0054	0.0168	0.0422	0.0838	0.1199	0.0573	-0.018	-0.083	-0.24	-0.524	-0.818	-0.97	-0.806		
	16	0	1E-06	2E-05	0.0001	0.0005	0.0019	0.0058	0.0141	0.0275	0.0346	-8E-04	-0.015	-0.063	-0.201	-0.497	-1.028	-1.833		

Figure 19.8. Products of expected payoffs for Ann when betting correctly and frequencies. Values are in 1/1000s

The sums of the remaining values is negative, about $-0.7/1000$. The values are not the expectations, since we have fewer cases left, so they become more likely. We would have to multiply the values in the table by $\frac{27}{32-k}$ to obtain the expected payoffs. So the expected payoff is $\frac{27}{17} \cdot 0.7/1000 \approx -1/1000$. That means the casino has an advantage if reshuffling is done that early. Ann loses on average 0.1 cent per round, even when she counts cards and plays optimally.

In the Excel sheet [CardCounting.xlsx](#) you can experiment with different maximum bet values and reshuffling times and see when Ann would win. It may be obvious but it follows from this sheet that the longer the dealer plays until reshuffling, the better for Ann, and the larger the maximum bet the better for Ann. Here

are a few values:

- With maximum bet of 5, Ann has positive expectations if the deck is reshuffled when 7 or fewer cards are left.
- With maximum bet of 10, Ann has positive expectations if the deck is reshuffled when 11 or fewer cards are left.
- With maximum bet of 20, Ann has positive expectations if the deck is reshuffled when 14 or fewer cards are left.

Home Activity Let an optimal playing Ann play 10000 rounds against dealer Beth in the applet [MiniBlackJack3b](#). It will take some time, but the applet can run in the background. The first table below the game will be filled with the frequencies that the parameters occur at the start of a round. Compare the table with the table in Figure 19.5. After rescaling, both may look close, but there are some baffling discrepancies, for instance in the lower right corner. Try to explain this.

Home Activity Use again the applet [MiniBlackJack3b](#) to check whether Ann can win in the long run if the maximum bet is 10 and reshuffling is done when only 11 cards are left.

Exercises

1. In MINI BLACKJACK (14, 16), how likely is a win for Beth in position 2.-11 if Beth asks for another card?
2. In MINI BLACKJACK (11, 16), how likely is a win for Beth in position 11.-11 if Beth asks for another card?
3. In a game with five rounds, your winning probabilities are 0.2, 0.3, 0.7, 0.3, and 0.6. In each round you bet an amount from the minimum bet 1 to the maximum bet b . If you win a round, you get the bet back, plus the betting amount as a win. If you lose a round, your bet is gone. How would you play, and for which value of b is the game fair?
4. In a game with five rounds, your winning probabilities are 0.2, 0.3, 0.4, 0.5, and 0.4. In each round you bet an amount from the minimum bet 1 to the maximum bet b . If you win a round, you get the bet back, plus the betting amount as a win. If you lose a round, your bet is gone. How would you play, and for which value of b is the game fair?
5. Explain when you would bet high in real blackjack. How high?
6. In real blackjack, would shuffling earlier increase or decrease the player's expectation if the player is playing optimally and counting cards? Explain why.

Project 32

8 ROUNDS BLACK OR WHITE Eight black and eight white balls are in an urn. Eight rounds are played where you can bet that a black or white ball is drawn. If you are right, you get your bet back plus 0.9 times your bet. If you are wrong, you lose your bet. The minimum bet is \$1, and there is a maximum bet requirement. The balls chosen are not put back into the urn, so after the last drawing only eight balls are left in the urn.

Find out how large the maximum bet must be for the game to be profitable for the player.

Project 33

x ROUNDS RED, GREEN, OR BLUE Four red, six green, and seven blue balls are in an urn. A number x of rounds are played where you can bet on whether a red, green, or blue ball is drawn. If you choose the right color, you get your bet back plus

- 2.5 times your bet for of a red ball, and
- 1 times your bet for of a green or blue ball.

If you are wrong, you lose your bet. The minimum bet is \$1, and the maximum bet is \$5. The balls chosen are not put back into the urn.

Show that 2 ROUNDS RED, GREEN, OR BLUE is not profitable for the player. How many rounds are necessary? For which number x is x ROUNDS RED, GREEN, OR BLUE profitable for the player? Use the Excel sheet [ThreeColors.xlsx](#).

Project 34

Provide a more precise analysis of MINI BLACKJACK by taking into account the probabilities for a round starting with 32, 31, 30, ... cards. The probability of a round starting with 31 cards is 0 if we don't shuffle during rounds. You can determine the probabilities experimentally as relative frequencies by observing the deck sizes at the beginning of rounds, using the applet [MiniBlackJack3b](#). Then do the calculations using [CardCounting.xlsx](#).

CHAPTER 20

Example: Duel

Prerequisites: Chapters 8, 12, and 16.

In this chapter we will discuss some cruel-looking games that model a 19th Century duel. Since dueling is a male game, Ann’s and Beth’s brothers Adam and Bob play it.

20.1 One Bullet

DUEL(1, 1|0.1, 0.2, . . .) Two players, Adam and Bob, alternate moves. At each move, the player decides whether to shoot or to walk. The probability of hitting the other is 0.1 at the first move and increases by 0.1 after each walking step. Each player has only one bullet. That means that after one shot, the game is over. If he hits the other, he has won. If he misses, he has lost (the other, gentleman he is, will not bother to move closer and shoot). Winning has a payoff of 1, losing without being hurt of -1 , and losing by being hit of -2 .

The sum of the payoffs is zero if nobody gets hurt, but negative otherwise. Duel is not a game you want to play! By changing the payoffs we would be able to create other variants.

Student Activity Play ten rounds of the game in the applet [Duel1111](#) to learn how to play.

We have a sequential game with randomness. The extensive form, together with its backward induction analysis, is in Figure 20.1.




Figure 20.1. Backward induction analysis

Although the game is not symmetric, for the sequence of hitting probabilities the expectations for both players are equal, and both are negative.

Variants of the game, where the hitting probabilities (and maybe the payoffs) are different, can all be solved using backward induction. It may be tedious, but Excel can help. Thus all we need is a computer and Excel and someone to prepare the Excel sheet for us, or knowledge of how to create such an Excel sheet. But would a “Just wait half an hour and I will tell you the solution” be an appropriate answer to somebody asking us about a new variant, let’s say where the hitting probabilities grow like 0.1, 0.15, 0.25, 0.3, 0.4, just to give an example? We will see in the next subsection how to analyze all these variants together without computer help.

20.1.1 Analysis of One-bullet Variants with Increasing Probabilities without Computer Help

There are patterns, and Excel could help find them. We could analyze ten cases, and see if there is a pattern. But sometimes a clever idea can spare us work. We will see how asking the right questions and thinking will find some patterns without computer help.

1. Monotonicity In the solution, with low hitting probabilities both players decide to walk early and shoot later. This may be obvious, but does it hold for all variants? Should a player who would shoot at some distance ever decide to walk later? For one play, this is irrelevant since the player would never see the later situations anyway—he would have shot before and the game would be finished—but for the analysis it may be an important question. Let’s look at situation I where Adam, playing rationally and optimally, decides to shoot, as in the part of the tree shown in Figure 20.2.




Figure 20.2. Part of the game tree

The only assumption we make is that the probabilities increase monotonically for each player as the distance is reduced. That means we assume $p_1 < p_3$, but we don’t have to assume $p_1 < p_2$ or $p_2 < p_3$.

Assume optimally-playing Bob would decide to walk when in situation II. But this would imply that Adam, by shooting, could expect at least $2p_3 - 1$ in situation III, and therefore also in situation II. But then he would not have chosen to shoot in situation I, since $p_1 < p_3$, a contradiction to our assumption. Therefore Bob would shoot in situation II. With the same reasoning, this would imply that Adam would shoot in situation III, and so on. If there is a distance where one or both would fire, then both would fire at every later time, at every shorter distance.

2. When to start shooting We look at the situation in the extensive form where one player, here Bob, shoots. In which case would Adam shoot in the step before? Let p_1 and p_2 denote the hitting probabilities. Then Adam has an expectation of $p_1 \cdot 1 + (1 - p_1) \cdot (-1) = 2p_1 - 1$ when he shoots. When he waits, Bob shoots, and Adam’s expectation in that case is $p_2 \cdot (-2) + (1 - p_2) \cdot 1 = 1 - 3p_2$. Therefore Adam will wait provided $2p_1 - 1 \leq 1 - 3p_2$, or $2p_1 + 3p_2 \leq 2$, and shoot otherwise. This also holds for Bob.

If the probabilities grow slowly and are the same for both players, so they have equal skills, then p_1 and p_2 are almost identical and the players would wait until the probability reaches about $2/5 = 0.4$.




Figure 20.3. Part of the game tree

20.1.2 Analysis of DUEL(1, 1|m, 2m, 3m, ...)

We assume constant increments for the probabilities. In the one-shot game DUEL(1, 1|m, 2m, 3m, ...) with fixed m , if Adam moves first, his hitting probability at his k th choice would be $a_k = (2k - 1)m$, and Bob's hitting probability at his k th choice would be $b_k = 2km$.

- When Adam decides, his hitting probability is $p_1 = (2k - 1)m$ and Bob's hitting probability in the following move is $p_2 = 2km$, for some k . Adam will walk if $2p_1 + 3p_2 \leq 2$, that is, if $2(2k - 1)m + 3(2km) \leq 2$, or $(4k - 2 + 6k)m \leq 2$, or $k \leq \frac{m+1}{5m}$.
- When Bob decides, we have Bob's hitting probability $p_1 = 2km$ and Adam's hitting probability in the following move $p_2 = (2k+1)m$. Therefore Bob will walk if $2p_1 + 3p_2 \leq 2$, i.e. if $2 \cdot 2km + 3(2k+1)m \leq 2$, or $(4k + 6k + 3)m \leq 2$, or $k \leq \frac{2-3m}{10m}$.

So who will shoot first? It depends on the numbers $\frac{m+1}{5m}$ and $\frac{2-3m}{10m}$. More precisely, it depends on the ceilings of these numbers, denoted by $\lceil \frac{m+1}{5m} \rceil$ and $\lceil \frac{2-3m}{10m} \rceil$, which are obtained by rounding up. If the numbers are equal, or if $\lceil \frac{m+1}{5m} \rceil \leq \lceil \frac{2-3m}{10m} \rceil$, then Adam shoots; otherwise Bob does. Nobody shoots second in the 1-bullet version, as explained above.

By the design of the game, the expectations for Adam and Bob are those in the situation where one starts shooting. If Adam shoots at move k , Adam's expected payoff is $2(2k - 1)m - 1$ and Bob's expected payoff is $1 - 3(2k - 1)m$. If Bob shoots first, at his k th move, Adam's expected payoff is $1 - 3(2km) = 1 - 6km$ and Bob's expected payoff is $2(2km) - 1 = 4km - 1$.

For example, take $m = 0.07$. We compute $\frac{m+1}{5m} = 3.06$ and $\frac{2-3m}{10m} = 2.56$. Their ceilings are 4 and 3, which implies that Bob shoots first, in his third move. Adam's expected payoff is $1 - 6km = 1 - 18 \cdot 0.07 = -0.26$, and Bob's expected payoff is $4km - 1 = 12 \cdot 0.07 - 1 = -0.16$.

Student Activity Play ten rounds of DUEL(1, 1|0.05, 0.1, ...) (we have $m = 0.05$ here) against the computer (playing optimally) in the applet [Duel13](#).

Modeling Note Equivalent Models or the Elusiveness of Mathematics Is it ethical to investigate features like duels, or military games of which we might not approve? Let us define the game "I LOVE YOU". Man and Woman want to say "I love you" to the other, since they love each other, and each knows that the other enjoys hearing it. But each one wants to be the first saying it—we all know how lame a muttered "I love you too" is. Coming home from work and shouting "Darling, I love you" may not be

convincing. It looks staged. You have to wait for the right moment. If we assume that the probability of success of a whispered “I love you” increases over time, we may have a game that essentially looks like the one-bullet duel. The same mathematical structure can have completely different meanings, as different as shooting a bullet into someone or saying “I love you”.

Another example is market entry. Two companies are working on a new product. They know that whoever releases it first may have success and see the other company fail. In the case of success, the product will gain the whole market rather quickly, and the other company will not be able to put their new product in the market. Thus the company that successfully releases their product first wins. But if the product fails, the company cannot release the product again. In that case, the other company can wait as long as necessary, and then release the perfect product, which will be accepted. Thus in the second case the company that did not release the product will win. The probability for success will rise over time, since the product gets better and better.

Modeling Note Utility Payoff based on other Utilities? You may have observed an inconsistency in the “I LOVE YOU” model: Being aware that your partner wants to be the first to say “I love you”, wouldn’t you wait so as to satisfy him or her? But wouldn’t he or she feel the same way? Basing one player’s utility on another’s leads to self-reference paradoxes and is usually avoided. We could take this into account and change the payoffs.

20.2 Two or more Bullets

DUEL(2, 2|0.1, 0.2, ...) Two players, Adam and Bob, alternate moves. At each move, the player to move decides whether to shoot or to walk. The probability of hitting the other increases from 0.1 at the first move, to 0.2, 0.3, and so on, but it increases only after one has walked. Each player has two bullets. The game is over if one player has been hurt, or if one player has used both bullets without hitting the other—then this player has lost. Winning has a payoff of 1, losing without being hurt of -1 , and losing by being hit of -2 .

Student Activity Play ten rounds of the game in the applet [Duel1221](#) to learn how to play.

The game seems more complicated than the 1-bullet version. If a player shoots, only in case of a hit do we get an outcome with payoffs attached. In case of a miss, the game continues, but with a new bullet distribution—2 versus 1 or 1 versus 2.

Therefore, in order to analyze game DUEL(2, 2|0.1, 0.2, ...), or more generally an r -bullets versus s -bullets game like DUEL($r, s|0.1, 0.2, \dots$), we need to define many other positions. Let $A(r, s, p, \dots)$ or $B(r, s, p, \dots)$ denote the position where Adam has r bullets, Bob has s bullets, the present hitting probability is p (with the understanding that the probability would grow if the shot does not hit and the game continues). The only difference is that Adam is about to move next in $A(r, s, p, \dots)$, and Bob in $B(r, s, p, \dots)$. All positions with the same numbers r and s can be drawn in one figure, as in Figure 20.4, and are connected by arcs. In the case of a shot that misses, we leave the figure and enter another layer of r versus $s - 1$ bullets or $r - 1$ versus s bullets, depending on who shot. Only for $r = 1$ or $s = 1$ we would have payoffs attached in




Figure 20.4. One layer of the game digraph for a multiple bullets version

case of a shot that misses. The whole game digraph consists of all the layers, all looking like the one in Figure 20.4, with exits from missing shot positions into other (lower) layers. The whole game tree is huge.

At position $A(r, s, p)$, either Adam takes a step and we arrive at $B(r, s, p + 0.1)$, or Adam shoots in which case we arrive at a random move position. We could say that we arrive either at the end position carrying a payoff of 1 for Adam and -2 for Bob (Bob is hurt) with probability p , or at position $B(r - 1, s, p)$ with probability $(1 - p)$. If $r = 1$, the second position is an end position—Adam has lost and gets a payoff of -1 , and the payoff for Bob is 1. Then backward induction can be performed by computing expected values first for positions $A(1, 1, p)$ and $B(1, 1, p)$ (see the preceding section), and for positions $A(1, 2, p)$ and $B(1, 2, p)$, for $A(1, 2, p)$ and $B(1, 2, p)$, and finally for $A(2, 2, p)$ and $B(2, 2, p)$. Even if we are not interested in the game where one player has one bullet and the other has two, we have to deal with this kind of game (assigning expected values to the positions $A(1, 2, p)$ and so on) to be able to solve the 2 and 2 bullet version. This is done in the Excel sheet [Duel.xlsx](#).

Student Activity Play ten rounds of this game in the applet [Duel4](#).

20.2.1 A few Cases of DUEL(2, 2|m, 2m, 3m, ...)

Here are a few cases, analyzed using the Excel sheet, for different probability increments m :

m	Adam	Bob	Adam	Bob	Adam	Bob	Adam's payoff	Bob's payoff
0.05	step	step	step	shoot	step	shoot	-0.2	-0.2
0.055	step	step	shoot	step	shoot		-0.3026	-0.0461
0.06	step	step	shoot	step	shoot		-0.2464	-0.1304
0.07	step	step	shoot	step	shoot		-0.1376	-0.2936
0.073	step	step	shoot	shoot	shoot		-0.2878	-0.325
0.08	step	shoot	step	shoot			-0.0848	-0.2768

Looking at the cases, we can make a few observations:

It seems that monotonicity discussed above also holds for the 2-bullet game. Players walk until they have reached a certain probability, which now seems to lie somewhere between 0.16 and 0.18. With hitting probabilities higher than that, they would always shoot provided both have still both bullets. However, if one of the players has only one bullet left and the other still has both, both wait a little longer. Surprisingly the

one with the two bullets waits still a little longer. Even though the player can spare one bullet, he still will wait a little longer.

Exercises

1. Can you tell, maybe even without performing a backward induction analysis, how the players would play if the hitting probabilities are constant, as in $\text{DUEL}(1, 1|0.1, 0.1, \dots)$ or $\text{DUEL}(1, 1|0.2, 0.2, \dots)$. Justify your answer.
2. Analyze the game $\text{DUEL}(1, 1|0.1, 0.2, 0.3, 0.4, 0.3, 0.5, 0.6, 0.7, 0.8, 0.9)$, where the hitting probabilities are not monotonic.
3. Explain why knowing the solution of $\text{DUEL}(1, 1|0.1, 0.2, 0.3, 0.4, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9)$ and of $\text{DUEL}(1, 1|0.1, 0.2, 0.3, 0.3, 0.3, 0.5, 0.6, 0.7, 0.8, 0.9)$ imply that we also know how to play $\text{DUEL}(1, 1|0.1, 0.2, 0.3, 0.4, 0.3, 0.5, 0.6, 0.7, 0.8, 0.9)$ with non-monotonic hitting probabilities.
4. Analyze the game where both players decide simultaneously at which step in $\text{DUEL}(1, 1|0.1, 0.2, \dots)$ they will shoot. Either use a bimatrix, or justify your answer verbally.
5. All games we consider are finite in the sense that at every position only finitely many options are possible. Let's now discuss a simultaneous game, a continuous version of the 1-bullet duel game, with infinitely many options. Adam and Bob decide simultaneously the distances d_a and d_b between 1 and 0 at which they would shoot. Then the play is performed as follows. Both players start at a distance 1. Then it is decreased until the smaller of d_a and d_b —say it is d_a —is reached. Then the corresponding player—Adam under our assumption—shoots and the game is over. The hitting probability at distance d is $1 - d$. Winning has a payoff of 1, losing without being hit of -1 , and losing by being hit of -2 . Can you analyze the game and give some recommendation how to play?

Project 35

Drunk Adam

- a) How would the analysis of the one-shot sequential duel change if Adam's probability for hitting is, for any distance, half the probability for Bob hitting?
- b) What would change in the 2-bullet game?

Project 36

How more dangerous weapons affect the state budget and the health of citizens Assume the payoff of -2 when someone is hurt is calculated as the sum of -1 for losing and -1 for the person's share of the cost for doctors and recovery. Assume the state has expenses matching the share. The question is: If weapons become more dangerous, let's say doubling the costs, meaning that the payoff for the player when hurt decreases to -3 (-1 for losing and -2 for the costs), would the costs taken by the state (which are now twice the number of people hurt) also double, assuming that the total number of duels remain constant? Would the percentage of people being hurt remain the same? Investigate the sequential one bullet and the two-bullet models.

Project 37

Selecting m between 0.04 and 0.13

- a) Assume Adam can choose the step width m in DUEL(1, 1| m , 2 m , ...) or in DUEL(2, 2| m , 2 m , ...). Which m should he choose to maximize his expectations?
- b) Assume Bob can choose the step width m in DUEL(1, 1| m , 2 m , ...) or in DUEL(2, 2| m , 2 m , ...). Which m should he choose to maximize his expectations?
- c) Assume the state can choose the step width m in DUEL(1, 1| m , 2 m , ...) or in DUEL(2, 2| m , 2 m , ...). Which m should he choose to minimize the percentages of injuries?

Project 38

What duels are best for society? The state decides to allow one sort of dueling. How many bullets—3,3 or 2,2 or 1,1 and what probability increments should be recommended for the official duel so as to minimize injuries?

CHAPTER 21

Santa Monica in the 50s

“The RAND Corporation’s the boon of the world
They think all day long for a fee.
They sit and play games about going up in flames
For counters they use you and me.”

The RAND hymn
Malvina Reynolds, 1961.

In 1969, as a small German boy, I watched together with millions of others as Americans set foot on the moon. For me this was an incredible demonstration of the power of science and technology and of US leadership in those fields. It seemed that anything was possible.

I was too young to have heard about the threats of the cold war, about the treatment of the blacks in the US, and the resulting riots, about assassinations of presidents and Martin Luther King, about protests in 1968 in all the western world, and about the Vietnam war, although I must have seen some glimpses of that on my grandma’s newly acquired TV. Instead I drew my knowledge about the US from the *Big Book on America*, which I got from our local library. This book was from the 50s and painted a very optimistic picture of the US, with all people nice, friendly, and rational. Happy (white) families in nice, clean houses in the suburbs, the men driving to offices in huge cars that resembled ships, dressed in white short-sleeved shirts and dark suits, heading to interesting jobs in huge air-conditioned buildings. Nice, functional, modern buildings, like the RAND building in Santa Monica.

The RAND Corporation was founded in 1946 as a think tank as part of the Douglas Aircraft Company. In its earliest days about 200 researchers worked on their own projects or analyzed problems given them by the US Air Force. In 1948 Douglas spun RAND off into a separate corporation. The most recent technologies have always been of interest to the military, but during World War II the importance to the military of scientists and mathematicians grew dramatically. Mathematicians played crucial roles in developing the atomic bomb in the Manhattan project, in improving the first computers, and in researching and practicing encryption and decryption. After World War II the US and the Soviet Union began the cold war, and think tanks like RAND grew. Among the subjects useful to the military were operations research, in particular linear and dynamic programming, and network flows, and game theory.

In the 50s RAND became the second center of game theory, Princeton being the first. In those days, although he was just a consultant, RAND’s most famous participant was John von Neumann. Von Neumann was on the faculty at Princeton’s Institute for Advanced Studies, and was one of the world’s leading mathematicians. Many Princeton game theorists either visited RAND every summer (like John Nash) or went there permanently (like Lloyd Stowell Shapley, who worked at RAND from 1954 to 1981). George Dantzig was a RAND employee. He is considered to be the father of linear programming, an area closely related to zero-sum two-player games. George W. Brown was also a RAND employee, from 1948–1952, and worked in comput-

ers and in game theory, formulating in 1951 his fictitious play approach. Thomas Schelling had fifty years of affiliation with RAND, and Robert Aumann was a RAND consultant between 1962 and 1974. Schelling and Aumann shared the Nobel Memorial Prize in Economics in 2005. Also influential in advocating game theory was John David Williams, the head of the mathematics department. See [P1993] and [N1998] for descriptions of game theory at RAND.

RAND provided mathematicians with an ideal atmosphere for innovation, apparently giving them total freedom to choose projects. The theory of games developed substantially. But what about practical applications? RAND mathematicians created and investigated simple models of warfare and later of negotiations and arms control. RAND researcher Melvin Dresher's 1961 book on game theory [D1961] discusses many games in military terms, like tactical air-war games. Another example of an early model developed at RAND is the COLONEL BLOTTO game that exists in different variants (some discussed in Dresher's book). Here is a discrete version:

COLONEL BLOTTO(4, 9, 9) Two players have nine armies each that can be sent to four battlefields of about equal importance. The player sending more armies to a battlefield wins it. The player winning more battlefields wins the game.

Student Activity Play ten rounds of the game against the computer in the applet [BlottoC](#).

How did you like the game? Was it boring, with a lot of draws? Are you troubled with the idea of playing with tanks? Surely some of you have tried shooter games, and have seen much more violence on screen. Also a game is a game is a game, and playing the game is different from deploying real tanks. But if you are uneasy playing games with tanks, try the following more peaceful game, a variant of ELECTION discussed in Chapter 7 with four districts of equal importance.

ELECTION(4, 9, 9) Two candidates for an office have nine resources each that can be sent to four districts of about equal importance. The player sending more resources to a district wins it. The candidate winning more districts wins the election.

Student Activity Play ten rounds of the game against the computer in the applet [ElectionC](#).

The game has the same structure as COLONEL BLOTTO(4, 9, 9): the same choices, the same payoffs, the same rules for who wins. Mathematicians call structures like these that are identical except for names and labels **isomorphic**. So here we can make the standard response of mathematicians if somebody tries to make them responsible for applications of their findings: The essence of mathematics is abstraction, and there are usually many applications of the same abstract model—nice ones and ugly ones.

I still find it troubling that games like this one could really be used for war.

Some have criticized RAND mathematicians for the opposite reason. Some in the military doubted the usefulness of mathematical models in war. The 50s seem to have been filled with an almost total belief in mathematics, science, and technology, but it crumbled in the 60s [P1993]. Even for someone opposed to wars, games developed at RAND might look harmless, even silly. The name COLONEL BLOTTO may indicate that the game was not taken seriously by its inventors.

There are a few more objections to the application of game theory to war. Game theory requires few players, but millions can be involved in wars. Therefore models of war usually consider only decisions of the army leaders, generals and the government. Only they are the players. Game theory also assumes

rationality. Isn't war totally irrational? This depends on your point of view. If you look at a simple game of two countries of about the same power, and two options for each—declaring war or not, declaring war is not a Nash equilibrium option, since the probability for winning may not be much larger than 50%. But many people on both sides will die. In addition, win or lose, each side will suffer significant destruction of infrastructure. War is not a zero-sum game. But if one side can win easily and with high probability, the expected gain may be considered higher than the expected loss. You may also object to having human lives as part of the payoffs, together with gain and gain of territory, power, and influence. Phrases like “total expected kill rate” used in some of these papers may repel you. Isn't human life priceless? Yes, our life is priceless for us, but society attaches prices to human lives, and always has. Because of cost not every disease is treated with the most effective medicine, our cars and airplanes are not as safe as they could be, and so on.

A way to decrease the inclination towards war is for both sides to have powerful weapons, capable of great destruction. This is mutual deterrence. Work at RAND on mutual deterrence influenced a number of policies in the early years of the cold war, the threat of Mutually Assured Destruction (MAD), and the doctrine of massive retaliation—the threat to retaliate fiercely, maybe with nuclear weapons, to any attack. Threats and their credibility and the theory of military balance are at the heart of the idea to keep peace by mutual deterrence. Technological progress made changes in the models and policies necessary. Because of intercontinental ballistic missiles, measures had to be taken to preserve second strike ability. RAND was influential in the shift of focus from first strike possibilities to strengthening second strike capabilities.

The arms race can be formulated as a simple simultaneous game. Assume there is a new technology that needs some time to develop, and work on it can be hidden. Assume that it gives a country a huge advantage, so that a war could be won without much loss, provided the other country doesn't have the new technology. If both countries have it, it is a loss for both because of the huge costs of developing and building the weapons. Examples are the H-bomb, which was built by the Soviet Union and the USA in the early 50s, intercontinental ballistic missiles, and (much later, in the 80s) strategic defense initiatives, a ground and space-based protective shield against intercontinental ballistic missiles. In such cases countries play a simultaneous game with the options of developing the new technology or not. The payoffs make it a classical prisoner's dilemma game. By the way, around 1950, the prisoner's dilemma game was invented at RAND by Merrill Flood and Melvin Dresher and analyzed thoroughly there. RAND even ran experiments to see how non-game theorists would play.

Even those who were critical towards the arms race must admit that mutual deterrence succeeded to some degree. There was no war between NATO and Warsaw Pact members in all the years the Warsaw Pact existed (1955 to 1991). But this came at a high cost: in those years enormous amounts of money were spent on the military, and we lived with some fear about the terrible things that could happen.

In spite of the optimism and belief in rationality, technology, and progress, the 50s must also have been scary years, with all the threats: first the atomic bomb, then the H-bomb, then intercontinental ballistic missiles, all managed by unflinching and uncompromising Soviet and US leaders. During its teenage years, game theory spent a lot of time playing with such models of war in an atmosphere of angst and hostility. It might have tainted its character.

CHAPTER 22

Theory 6: Extensive Form of General Games

Student Activity Play VNM-Poker in the applet [VNMPokerseq13](#). The description can be found there.

Student Activity Play 2-round WAITING FOR MR. PERFECT in the applet [Waiting2b](#).

Chapter 2 was about simultaneous games. We showed how a matrix representation, the normal form, can be helpful in analyzing a game. In Chapters 8 and 16, which dealt with sequential games with or without randomness, we learned how to describe such games in extensive form, and how these games have a clearly defined solution and (expected) value, which we can compute using backward induction. An important assumption was that we had perfect information about the games.

However, most games are neither simultaneous nor purely perfect-information sequential. They are something in between. Some are essentially sequential with some moves done in parallel, like WAITING FOR MR PERFECT; some are sequential but lack perfect information like VNM-POKER. Perfect information means that every player is aware of all previous moves of all players. Simultaneous moves can be rephrased as sequential moves with non-perfect information—we put them into any order without giving the player who moves second knowledge about the move of the player moving first. In this chapter we will see that the description in extensive form can be applied to any game, even one with imperfect information.

22.1 Extensive Form and Information Sets

As with sequential games with perfect information, games with complete but possibly imperfect information can be described by a tree or a digraph, a so-called **extensive form**, as was first shown by Kuhn in 1953 [K1953]. We need one additional concept, that of information sets.

The idea is simple. We still have all of the vertices (positions), including a start position and end positions at which the payoffs of the players are given. Every non-end position belongs to one player or is a random position where a random move will occur. With non-perfect information a player to move may not know exactly at which position he or she is. For instance, cards can be dealt but the player doesn't know the cards in the other player's hands, or another player has moved but our player doesn't know how. The player to move cannot distinguish between several positions. The positions in which the game could be at that moment are combined into an **information set**. There are requirements they must satisfy:

- Information sets belong to players. Each information set of a player contains only positions of that player.
- Different information sets are disjoint.
- All vertices in the same information set have the same number of moves, and they have the same names for each vertex. Otherwise the information of which moves are possible would give the player a hint about the actual position.

- There are a few more requirements for information sets:
 - If a vertex A is an (iterated) successor of vertex B, then both must be in different information sets.
 - There are no time loops: if A2 is an (iterated) successor of A1, and B2 an (iterated) successor of B1, then A1 and B2 must be in different information sets, or A2 and B1 must be in different information sets.
 - We have perfect recall: All immediate predecessors of vertices in an information set are in the same information set.

All the conditions are automatically satisfied if we start with a concrete game and draw its extensive form. Only when we try to draw an extensive form without having a game in mind do we have to make sure that these requirements, and a few more, are met.

Let me again list the ingredients of the extensive form:

- We have an acyclic digraph consisting of vertices and directed arcs between them. **Acyclic** means without cycles: it is impossible to start at a vertex, move along arcs in the indicated direction, and return to it. Acyclic digraphs have representations where the arcs are directed from left to right, so the arcs can be and usually are suppressed in the representation. We have exactly one **start vertex** or **origin** with no incoming arcs. The arcs leaving a vertex are called the **alternatives**, **options**, or **moves** at that vertex. The vertices with no outgoing arcs are called **final,terminal**, or **end vertices**. If we have information sets with more than one vertex (i.e., imperfect information), then, as we will explain, we might want to use a (directed) tree, a **game tree**, instead of a general digraph to represent the game.
- All vertices except the final vertices belong to exactly one player or are random vertices, where the random moves occur.
- Payoffs for players are attached to final vertices.
- At every random vertex all outgoing arcs have labels giving their probabilities, and the numbers add up to 1.
- One or several vertices belonging to the same player are combined in information sets. Every vertex belongs to an information set, possibly comprising just one vertex. Different information sets do not overlap. Vertices in the same information set look alike insofar as they have the same number of outgoing arcs (alternatives) labeled in a consistent way.

We will now look at some examples.

Example 1 Assume Ann and Beth play a two-player simultaneous game with two moves each. The general normal form is:

	B_1	B_2
A_1	$A_{1,1}, B_{1,1}$	$A_{1,2}, B_{1,2}$
A_2	$A_{2,1}, B_{2,1}$	$A_{2,2}, B_{2,2}$

Then, as explained, we remodel the game as a sequential game, where Ann decides first, and Beth, without knowing Ann’s decision, decides next. Then Beth has an information set with two positions. When deciding her move, she does not know which she is in. The extensive form is shown in Figure 22.1.




Figure 22.1. Extensive form of a simultaneous game

Example 2 PRISONER’S DILEMMA OR STAG HUNT There are two players who first simultaneously vote whether they want to play PRISONER’S DILEMMA or STAG HUNT. If both vote for the same game, it is played. Otherwise no game is played, and both players have a payoff of 0. The two games are simultaneous and have the following bimatrix normal forms:

	Cooperate	Defect
Cooperate	2, 2	0, 3
Defect	3, 0	1, 1

	Hunt stag	chase hare
Hunt stag	3, 3	0, 2
Chase hare	2, 0	1, 1

The extensive form is shown in Figure 22.2




Figure 22.2. PRISONER’S DILEMMA OR STAG HUNT

Example 3 LE HER*(2, 4) Ann and Beth play with four cards of value 1 and four cards of value 2. Each randomly draws a card and looks at it without showing it to her opponent. If Ann holds a 1, she can exchange cards with Beth (without knowing what card Beth holds). If Beth now holds a card of value 1, she has the opportunity to exchange her card with a randomly chosen card from the deck. Then both players reveal their cards, and the player having the higher card wins one dollar from the other.

Versions of this game, variants of an old game, have been introduced by Karlin ([K1959], p. 100), and have been discussed and analyzed in [BG2002]. In the versions discussed in the literature, both players can also swap cards if they hold a card of highest value, but why would they do this? So we may as well forbid swapping cards of the highest value (in our case of value 2), as we did in our description of LE HER*(2, 4). This keeps the extensive form small and manageable. Figure 22.3 shows the extensive form of LE HER*(2, 4).

The game starts with the random move of dealing the hands. The pair (1, 1) (a 1 to both Ann and Beth) is dealt with probability $\frac{3}{14}$, the same probability as for a (2, 2), whereas the combinations (1, 2) and (2, 1) each have probability $\frac{4}{14}$. Next it is Ann’s turn. There are four positions, but Ann, being unaware




Figure 22.5. Third extensive form of LE HER*(2, 4)

In Figure 22.5 we get another extensive form by playing the game in a different order. First Ann gets her card and decides whether she wants to exchange it. After that Beth gets her card and decide on her move. From a mathematical point of view the game is identical to the one described in Figure 22.4. The extensive form could be reduced by not accepting one-option moves as moves. See Exercise 1.

22.2 No Backward Induction for Imperfect Information

Backward induction can be tried in any sequential game, but it is stuck at information sets containing at least two vertices. In the example in Figure 22.6 we can and did assign values to Ann’s decision between A3 and A4, to Ann’s decision between A5 and A6, and to the random move. But we are stuck at Beth’s move. If Beth decides B1, her payoff could be 3 or 0. If she decides B2, her payoff could be 1 or 2. It is difficult to tell what Beth should do and what she would expect, especially since she cannot tell which position is more likely.




Figure 22.6. A game with imperfect information

22.3 Subgames

Some extensive forms of games can be broken into two parts along a vertex where there is almost no connection between the remaining parts. The part, let's call it the branch that follows after the vertex where we cut now forms a start vertex and can be played without reference to the other part. The preceding part, which we call the trunk, and for which the cut vertex now forms an end vertex, can be played without reference to the branch, except that this newly created end vertex has no payoffs assigned yet. If the branch game has clearly defined expected values, which is the case, for example, if it has exactly one pure Nash equilibrium, it would be reasonable to assign the expected values as payoffs for the vertex in the trunk.

Breaking a game into two smaller games called **subgames**, may simplify the analysis, since they have less complexity than the original game. In sequential games with perfect information every vertex could serve as a cut vertex. If we have a sequential game with imperfect information, or a nonsequential game reformulated as one, then in addition to the arcs in the game tree, there are the information sets, which form another type of connection between vertices. When we try to cut a game into two parts, we have to take care that there is no connection between them except at the cut vertex.

In the PRISONER'S DILEMMA OR STAG HUNT example, there are only two vertices at which we could cut the game into two parts: Ann's moves that are not the start move. In the example both vertices are the only non-start non-end vertices with a one-vertex information set. We will see later that such a vertex x that forms its own information set does not always allow for cutting the game at x . We need a second condition, namely that for every successor vertex y of x , all vertices in the same information set as y must also be successors of x .

We call such a possible cut vertex, together with all its successors, a **branch subgame**. Thus a branch subgame consists of a vertex x and all its successors if

- x is a single-vertex information set, and
- all information sets are totally inside, or totally outside, the successor set of x .

For our PRISONER'S DILEMMA OR STAG HUNT example, let's look at the PRISONER'S DILEMMA subgame in Figure 22.7. It has one pure Nash equilibrium, defecting versus defecting, leading to payoffs of 1 to both players. Since we know that a single pure Nash equilibrium is likely to be chosen for rational play, in the overall game we could stop when both players have decided to play the PRISONER'S DILEMMA subgame, since we know what the outcome will be. So instead of looking at the overall game, we can look at the truncated game in Figure 22.8 with payoff of 1 and 1 at the (now terminal) vertex where the subgame has been cut off. We can do this because we know that both players will get a payoff of 1.

For the other subgame, STAG HUNT, things are more complicated to agree on expected values Z for Ann and W for Beth in the subgame. The reason is that we have two pure Nash equilibria—stag versus stag with a payoff of 3 for each player, or hare versus hare with a payoff of 1 for each. Though we might think that $Z = 3$, $W = 3$ are the right values for this subgame, $Z = 1$ and $W = 1$ is also possible. Then, if we cut off the subgame and introduce payoffs 3,3 or 1,1 at the (now terminal) cut vertex, we arrive at the totally truncated game in Figure 22.10.

The game, which is the first simultaneous round of our two-round game, has two Nash equilibria: both players select the same game for the second round. We don't know what payoffs to expect, 1 and 1, or 3 and 3. The uncertainty of the subgame affects the overall game.

Let's demonstrate the method for our LE HER*(2, 4) example, on the first extensive form in Figure 22.3. We can identify four small subgames, starting at the random move after Beth's swap, and two larger subgames starting at Beth's single-vertex information sets after Ann has swapped. Two of the smaller subgames are also subgames of the larger ones, so in what follows we will focus on the remaining four subgames. All are perfect-information sequential games, so they can be analyzed using backward induction. If we cut off the four branches and insert the backward induction values of the subgames at the new terminal vertices (indicated

by the white diamonds) we get the truncated extensive form in Figure 22.11. If we do the same for the smaller extensive form in Figure 22.4 where the moves with one options have been removed, we arrive at the truncated extensive form in Figure 22.12.

22.4 Multi-round Games

Many games are played in rounds, in which players move simultaneously. After each round the simultaneous moves may be followed by random moves. You played an example at the beginning of the chapter: WAITING FOR MR. PERFECT. We can reformulate such a game as a sequential game with nonperfect information, where in each round Ann moves first, followed by Beth and so on, with either of them not knowing how the other has moved until the round is finished. Then all information sets for Ann are single-position information sets, and they all start subgames. The method of chopping off subgames can accordingly be used to analyze such a game.

22.5 Why Trees for Imperfect Information?

So far we have used trees and digraphs for extensive forms. For a sequential game with perfect information we should identify positions belonging to the same player that have identical futures (including all payoffs at ensuing end vertices).

This is not true for games with imperfect information. Nontrivial information sets (containing each at least two vertices) with identical futures can no longer be identified. Thus, for imperfect information, we need to use trees.

In the example of Figure 22.13 there are two information sets for Ann with identical futures but different pasts. Although the futures are identical, the optimal player would play differently in them. Why? Since Ann has learned something from Beth’s behavior? Yes. Ann knows something about her own position, at which vertex in the information set Ann is. If Ann has the first information set, she can be quite sure that Beth did not move before, since Beth would likely have moved B_1 , assuring her a decent payoff of 1. If Ann has the second information set, two positions are possible.

We still can identify one-vertex information sets with identical futures in games with imperfect information. So in special cases digraphs can be used, but with caution.




Figure 22.8. The trunk




Figure 22.10. The trunk

Exercises

1. Reduce the extensive form given in Figure 22.5 by eliminating moves with only one option.

2. Draw the extensive form of the following game:

Both Ann and Beth put one dollar in the pot. Ann gets a card from a stack of four queens and four kings and looks at it privately. Then Ann either folds, in which case Beth gets the money in the pot, or raises. Raising means that Ann has to put another dollar in the pot. When Ann has raised, Beth either folds, in which case Ann gets the pot, or Beth calls by putting also one more dollar in the pot. If Beth calls, Ann gets the pot if she has a king, otherwise Beth gets the pot.

3. Draw the extensive form of the following game:

Both Ann and Beth put one dollar in the pot. Ann gets a card and looks at it privately. Then Ann either checks, in which case Ann gets the money in the pot if Ann's card is red, or Beth gets the pot if Ann's card is black. Ann can also raise by putting another dollar in the pot. Now Beth either folds, in which case Ann gets the pot, or Beth calls by putting one more dollar in the pot. If Beth calls, Ann gets the pot if she has a red card, otherwise Beth gets the pot.

4. Draw the extensive form of the following game:

Ann starts the game by selecting (Ann doesn't draw, she chooses) two cards from a deck of cards containing four queens and four kings. Ann puts them face down in front of Beth. Beth is allowed to see one




Figure 22.11. Fourth, pruned extensive form of LE HER*(2, 4)




Figure 22.12. Fifth, pruned extensive form of LE HER*(2, 4)




Figure 22.13. A game

of them. Then Beth must guess whether the two cards are two kings, two queens, or one king and one queen. If she is right, she wins \$1 from Ann, otherwise she has to pay \$1 to Ann.

- 5. Draw the extensive form of the following game:
Ann starts the game by selecting (Ann doesn't draw, she chooses) one card from a deck of cards containing four queens and four kings. Ann puts it face down in front of Beth. Beth is not allowed to see it, but is allowed to see one card of the remaining deck of seven cards. Then Beth must guess whether the card face down is a king or a queen. If she is right, she wins \$1 from Ann, otherwise she has to pay Ann \$1.
- 6. Draw the extensive form of the following game:
Ann starts the game by selecting (Ann doesn't draw, she chooses) two cards from a deck of card containing four queens and four kings. Ann puts them face down in front of Beth. Beth is not allowed to see them, but is allowed to see one card of the remaining deck. Then Beth must guess whether the two cards

are two kings, two queens, or one king and one queen. If she is right, she wins \$1 from Ann, otherwise she has to pay Ann \$1.

CHAPTER 23

Example: Shubik Auction II

Prerequisites: Chapters 8, 12, 16, 22, and 10.

In this chapter we look at this simultaneous game with randomness, and we discuss connections to games with nonperfect and incomplete information. This is a continuation of Chapter 10, where we saw that knowing in advance the maximum number of moves results in a disappointing optimal solution, where the player who will not have the last move will not even start bidding. What happens if the number of bidding rounds is finite but unknown? Or if the number of rounds is finite, but after every move the game could randomly end?

23.1 Possible Sudden End

In SHUBIK AUCTION, the player with the last move will bid and the other will pass immediately. What happens if we don't know in advance which player has the last move? Assume there is a maximum number of rounds, and assume that the game can terminate after each round with probability p . This makes the game fairer, more interesting, and, as we will see, more profitable for the auctioneer.

SHUBIK AUCTION(A, B, n, p) Two players, Ann and Beth, bid sequentially for an item, with bids increasing by increments of \$10. The item has a value of A for Ann and B for Beth. The game ends if one player passes, i.e., fails to increase the bid, or after Ann and Beth complete the n th bidding round. There is a third way the game could end: after every bid, the game may terminate with probability p . After the game ends, both players pay their highest bids, but only the player with higher final bid gets the item.

For $p = 0$ we get the non-random version discussed in Chapter 10. How does the backward induction analysis change when $p > 0$?

Let us look at the example SHUBIK AUCTION(35, 35, 6, 0.5). The game tree is in Figure 23.1, which shows the expected payoffs for all positions and the recommended moves obtained from backward induction. In Beth's last position, she will bet 60. That implies that the expectations at this position for Ann and Beth are -50 and $35 - 60 = -25$. If the game ends before Beth can bet 60, then Ann will receive the item and get a payoff of $35 - 50 = -15$, and Beth's payoff will be -40 . Therefore the expectations at the last random position are the averages of $(-50, -25)$ and $(-15, -40)$, namely $(-32.5, -32.5)$. Since -30 is larger than -32.5 , Ann passes instead of bidding 50 in the move just before that random position. Proceeding this way, we get expected payoffs of 7.5 for both players at the beginning of the game.

Ann plays differently than when there can be no random ending. As there, she will not bid except at her first move. Why the change? There is a large probability that the game will end after the first bid, and Ann will get the item when bidding is still cheap.




Figure 23.1. Game tree for SHUBIK AUCTION(35, 35, 6, 0.5)

Is it possible to do the analysis once and for all for arbitrary parameters A , B , and p ? For fixed $n = 6$, the structure of the game tree remains the same, and the payoffs can be expressed in terms of A and B , as shown in Figure 23.2. Ann's positions are labeled as A1, A2, A3, Beth's positions as B1, B2, B3, and the random move positions as R1, R2, ..., R5.




Figure 23.2. Game tree for SHUBIK AUCTION(A , B , 6, p)

- In position B3, Beth has payoffs of $B - 60$ versus -40 , so she will bid 60 provided $B - 60 > -40$, i.e., if $B > 20$.
- The expected payoffs for Ann and Beth when at R5 are $(1 - p)(-50) + p(A - 50) = pA - 50$ for Ann and $(1 - p)(B - 60) + p(-40) = (1 - p)B + 20p - 60$ for Beth if $B > 20$. If $B \leq 20$, they are $A - 50$ and -40 .
- At A3, if $B > 20$, Ann has an expected payoff of $pA - 50$ and a payoff of -30 . The first number is larger if $A > 20/p$; then Ann will bid. If $B \leq 20$, Ann bids if $A > 20$. Accordingly Ann's and Beth's expected payoffs at A3 are

$$\begin{array}{ll}
 pA - 50 \text{ and } (1 - p)B + 20p - 60 & \text{if } B > 20 \text{ and } A > 20/p, \\
 -30 \text{ and } B - 40 & \text{if } B > 20 \text{ and } A < 20/p, \\
 A - 50 \text{ and } -40 & \text{if } B \leq 20 \text{ and } A > 20, \\
 -30 \text{ and } B - 40 & \text{if } B \leq 20 \text{ and } A < 20,
 \end{array}$$

If $B > 20$ and $A = 20/p$ we get a tie, and it is not clear what Beth's expected payoff would be. So we omit this case, as well as the case $B = 20$ and $A = 20$.

- The second and fourth cases could be combined. Thus, the expected payoffs in position R4 for Ann and Beth are

$$\begin{aligned}
 & (1-p)(pA-50)-30p \text{ and } (1-p)((1-p)B+20p-60)+p(B-40) \\
 & \hspace{15em} \text{if } B > 20 \text{ and } A > 20/p, \\
 & -30(1-p)-30p \text{ and } (1-p)(B-40)+p(B-40) \\
 & \hspace{15em} \text{if either } (B > 20 \text{ and } A < 20/p) \text{ or } (B \leq 20 \text{ and } A < 20), \\
 & (1-p)(A-50)-30p \text{ and } -40(1-p)+p(B-40) \quad \text{if } B \leq 20 \text{ and } A > 20.
 \end{aligned}$$

Are you still with me? If so, you can see that the analysis can be done with arbitrary parameters. Unfortunately the number of cases increases at each of Ann's and Beth's positions. Instead of pursuing this further, we will take a different approach, making numerical computations for different values A , B , n , p , and collecting the results in tables. For this, we use the Excel sheet [ShubikAuction.xlsx](#). There are different sheets for different values of n . On each sheet, when you input A , B , and p , the expected payoffs for Ann and Beth are displayed below each vertex.

For instance, in the 6-round sheet input $A = 35$ and $B = 35$. We get the chart in Figure 23.3. The lines jump at some values of p , around $p = 0.27$ and around $p = 0.57$. These are the values where Ann's or Beth's optimal strategy changes. For p between 0 and 0.28 Ann will always pass and Beth will always raise the bid. For p between 0.28 and 0.57 we get the same bidding pattern as discussed above for $p = 0.5$: Ann bids 10, but doesn't try higher bids, whereas Beth always bids. For larger p , both will always bid. The probability for a sudden end is just high enough to make bidding worthwhile. Ann's strategy changes around $p = 0.28$ and around $p = 0.57$, but it is Beth's payoff that dramatically changes at these values of p .




Figure 23.3. Payoffs for SHUBIK AUCTION(35, 35, 6, p)

If Ann values the item more, say $A = 45$, she will try harder to get it. We get the chart in Figure 23.4, which has four strategy patterns. For p between 0.45 and 0.55, Ann always bids, whereas Beth would pass on her 20 bid (while being still prepared to bid 40 and 60, but, these positions will not be reached). This strange behavior has to do with Ann bidding for smaller p , due to the higher worth of the item for Ann, whereas this mid-level value of p generates an inefficiently large payoff for Beth, given its low worth for her.




Figure 23.4. Payoffs for SHUBIK AUCTION(45, 35, 6, p)

23.2 Imperfect and Incomplete Information

Let us play a different variant now, a variant where instead of having after every move a random decision whether to terminate the game, when to terminate is decided randomly at the beginning but not revealed. Thus we have a maximum number of n moves, and the game starts with a random move that decides the actual number of moves that will be played. Assume that the number will be 1 with probability p_1 , 2 with probability p_2 , and so on, with $p_1 + p_2 + \dots + p_n = 1$.

Since Ann and Beth don't know how many moves they will play, the game is sequential with randomness and imperfect information. For appropriate parameters $p_1, p_2, \dots, p_n$ the game is just our SHUBIK AUCTION(A, B, n, p). Since the probability that SHUBIK AUCTION(A, B, n, p) ends after the first round is p , that it ends after the second move is $(1-p)p$, and so on, can select $p_1 = p, p_2 = (1-p)p, \dots, p_{n-1} = (1-p)^{n-2}p$, and $p_n = (1-p)^{n-1}p$. Figure 23.6 is the extensive form of the game equivalent to SHUBIK AUCTION(35, 35, 6, 0.5).

We can even describe SHUBIK AUCTION(A, B, n, p) as a game of incomplete information. We play at most n rounds, but it could be fewer. Incomplete information means that the extensive form is not fully known to the players.

In games of incomplete information players need beliefs about its structure—without them, players would not be able to make any decisions. If we assume that both players have the same beliefs about how likely it is that the length of the game will be 1, 2, $\dots, n$ moves, then we arrive exactly at the game with imperfect information discussed above. If the players have different beliefs, we could not model the game in this way.

23.3 The Auctioneer Enters the Game (optional)

Assume we play the 2-player variant, SHUBIK AUCTION(A, B, n, p), but before the start, the auctioneer has a move: deciding the value of p . Assume that the auctioned item is worthless to the auctioneer, so his or her payoff is the total amount of money paid. We arrive at a sequential 3-player game, which we call SHUBIK AUCTION w.A.(A, B, n) and in which the auctioneer has only one move, the first. How should the auctioneer place a value on p , and what payoff could he or she expect?

It depends on A and B , and probably to a lesser extent on n . But p is no longer a parameter of the game.

As can be seen from the chart in Figure 23.5, for $A = 30, B = 40$, and $n = 15$, the auctioneer should choose p slightly larger than $1/2$ or slightly larger than $2/3$ to obtain an expected payoff of almost \$15.




Figure 23.5. Payoffs for SHUBIK AUCTION w.A.(30, 40, 15)




Figure 23.6. SHUBIK AUCTION(35, 35, 6, 0.5) viewed as a game with imperfect information

Exercises

- Figure 23.7 shows the graph of SHUBIK AUCTION (45, 35, 14, p). Are the payoffs identical to SHUBIK AUCTION(45, 35, 6, p)? Explain.




Figure 23.7. Payoffs for SHUBIK AUCTION(45, 35, 14, p)

- Find the values p for SHUBIK AUCTION(35, 35, 6, p) where Ann's or Beth's payoffs "jump". Do the same for SHUBIK AUCTION(45, 35, 6, p) and SHUBIK AUCTION(45, 35, 14, p).
- Can you explain why, at the value p where Ann's strategy changes, Ann's payoff remains about the same but Beth's payoff jumps?
- Consider the version of the imperfect information variant discussed in Section 23.2, where there is a random and secret decision before the game starts whether to end after 1, 2, ..., 7 rounds, and where each option has probability 1/7. Formulate the game as one with perfect information, where after the first round the game ends with probability p_1 , after the second round with probability p_2 , and so on. How large are $p_1, p_2, \dots$?

Project 39

Consider the version of the imperfect information variant discussed in Section 20.2, where there is a random and secret decision before the game starts whether to end after 1, 2, ..., n rounds, and where each option has probability $1/n$. The game can also be formulated as a game with perfect information, where after the first round the game ends with probability p_1 , after the second round with probability p_2 , and so on. How do $p_1, p_2, \dots$ depend on n ? Modify the Excel sheet [ShubikAuction.xlsx](#) to model these games for various maximum game length. Explore a few cases.

Project 40

In SHUBIK AUCTION($A, B, 14, p$), where $20 \leq A \leq B$, will Beth ever pass?

Project 41

SHUBIK AUCTION(45, 35, 6, p) has four strategy patterns. Can you find A , B , n , and p where there are five or more strategy patterns?

Project 42

Three players in Shubik Auction with random end:

SHUBIK AUCTION(A , B , C , n , p) Three players, Ann, Beth, and Cindy, bid sequentially in cyclical order for an item worth A dollars to Ann, B dollars to Beth, and C dollars to Cindy. After every round the game may end with probability p .

Use the Excel sheet [ShubikAuction3.xlsx](#) to create the graph of payoffs of SHUBIK AUCTION(35, 45, 55, 16, p) for different p from 0 to 1. Do the same for two or three other games. Can you explain the findings? Do they make sense? Can you find a pattern?

CHAPTER 24

Theory 7: Normal Form and Strategies

24.1 Pure Strategies

In everyday language we distinguish between strategic and tactical behaviors. While tactical reasoning concerns what to do in a particular situation, the strategic point of view considers the whole picture and provides a long-term plan. In game theory, a strategy spans the longest possible time horizon—it is a recipe telling the player what to do in any possible situation. Since a situation translates into an information set in games with imperfect information, a pure strategy for a player lists all information sets in which the corresponding player has to move, together with rules on how to move in each information set. By choosing a pure strategy, the player decides on how to play in all possible situations (i.e., information sets). Even unlikely situations must be considered in advance.

In real life, few people start a game prepared with a pure strategy: they would start playing and decide what to do when it is their turn. So a pure strategy is more a theoretical than a practical concept. Other players or independent observers would not be able to decide whether a player plays with an a priori strategy. All they see is that decisions are made at different positions, but they do not know when they have been made. So we may as well assume that players have made all their decisions before the game starts.

Let's see how you can list, count, and encode the pure strategies a player has in a game. Since in principle in an information set every choice of move is possible, the product of the numbers of choices taken over all information sets of that player is the number of a player's pure strategies. We encode a player's strategies by first numbering all his or her information sets arbitrarily. We abbreviate each choice by a single letter. Then a pure strategy can be encoded as an n -letter word, where n is the number of the player's information sets, and the k th letter of the word tells what alternative the player chooses at the k th information set. If for example Ann has three information sets, the first and the last with two possible moves, say K, L and L, M, and the second with three possible moves, say M, O, and P, then Ann has the $2 \cdot 3 \cdot 2 = 12$ pure strategies encoded as KML, KMM, KOL, KOM, KPL, KPM, LML, LMM, LOL, LOM, LPL, and LPM.

For another example, recall that Ann has two information sets in the extensive form of LE HER*(2, 4) in Figure 22.3. She either has a card of value 1 or of value 2. In the first information set, she can either keep (K) or swap (S). In the second information set she must swap. Therefore Ann has $2 \cdot 1 = 2$ pure strategies KK, and SK. Beth has two information sets consisting of two vertices each, where Beth has a card of value 1 or 2 and Ann has kept her card. Beth also has two single-vertex information sets after Ann has swapped and Beth knows the values of both Ann's and Beth's cards. Except in the second information set, where Beth must keep, Beth can keep or swap. So Beth has $2 \cdot 1 \cdot 2 \cdot 2 = 8$ pure strategies, KKKK, KKKS, KKSK, KKSS, SKKK, SKKS, SKSK, and SKSS.

There were many extensive forms for this game. The smallest, in Figure 22.12, was obtained by neglecting moves with only one option, and cutting off perfect-information subgames. Here Ann and Beth have just one information set each, and both have just two pure strategies.

We have to ask the question, how could the same game have eight pure strategies for Beth in the first extensive form, and only two in the second one? Do our choices depend on the game’s description? Whether we list moves with only one option does not affect the number of pure strategies. When we eliminate subgames we assume that the player will choose the backward induction options in the subgames. This of course limits the pure strategies available.

24.1.1 Reduced Pure Strategies

When we look at the game whose extensive form is in Figure 24.1, we can see how problems may arise with pure strategies. This game, called 2-round chicken, is analyzed in in Chapter 33. Adam has two information sets, both single vertices. In each he has two choices, holding firm (H) or ducking out (D). So in principle Adam has four pure strategies HH, HD, DH, and DD. It doesn’t seem to matter whether Adam plays DH or DD: if he chooses D in the first information set, then he will never arrive at the second information set. So no matter what he planned to do, it will never happen. The strategy not to go to college and later accepting a professorship in mathematics at Harvard if offered, and the strategy not to go to college and later not accepting it if offered look different, and the second one looks impressive (“Look, I plan to decline a professorship at Harvard”), but both eventualities are the same.




Figure 24.1. An extensive form

We use the symbol “•” as a placeholder, meaning that any option available at the information set could be used (since in practice this information set will not occur when playing this strategy). In the example just considered D• would refer to both DD and DH. We say the resulting pure strategy is **reduced**. Pure strategies that cannot be identified with other pure strategies, like HH and also HD in the example are also called reduced. Therefore, Adam has three reduced pure strategies, namely HH, HD, and D•.

24.2 Normal Form

Assume all players start the game with pure strategies. Then the game’s outcome is determined if the players stick to their strategies. And why would they not? Strategies include how to react to surprising moves of opponents. Nothing should surprise a player playing with a fixed strategy, since the strategy tells what to do in every situation. So the players wouldn’t even have to play their game. All they would have to do to determine the game’s outcome would be to reveal their strategies simultaneously. Random moves would still have to be generated, so there might be different outcomes, but the uncertainty can be treated by looking at expected values. We can accordingly reduce complex games to simultaneous games with randomness, where players have their pure strategies as options. One advantage of this is that simultaneous games are easier to analyze. One disadvantage is that there may be a huge number of strategies.

For such a game’s **normal form**, also called its **strategic form**, we assume that we have a list of all pure strategies for all players. For each combination of pure strategies of the players, the payoffs, or expected payoffs in case the game involves randomness, are displayed. For two players, we get a bimatrix with two

entries in each cell, the payoff for Ann and the payoff for Beth. For three players we have a cube-like structure, or several of bimatrices, with cells containing three entries each, and so on.

The **reduced normal form** is defined similarly, except that only the reduced pure strategies are considered. In some cases the reduced normal form may be smaller than the normal form, but very often they have the same size.

For simultaneous games pure strategies and moves are the same, so the normal form is just their natural description. For purely sequential games of perfect information (with or without randomness) the solution found by backward induction yields a Nash equilibrium (in pure strategies) [K1953]. So for these types of games, there is no point in translating the extensive form into the normal form.

In our example LE HER*(2, 4), Ann has the pure strategies KK and SK, and Beth has the pure strategies KKKK, KKKS, KSKS, KKSS, SKKK, SKKS, SKSK, and SKSS in the extensive form in Figure 22.3. Let me show in one instance, SK versus SKSK, how to calculate the entries of the normal form bimatix. Since the game is zero-sum, we need calculate only Ann's payoff. We play the game and assume that both players stick to their chosen pure strategies. We have to consider all outcomes of the random moves. With probability $\frac{3}{14}$, both get a card of value 1. According to her pure strategy SK, Ann swaps and Beth then faces her third information set. Playing SKSK, Beth swaps, and the expected payoff for Ann is $\frac{-2}{3}$. With probability $\frac{4}{14}$, Ann gets a card of value 1 and Beth one of value 2. Ann swaps again; but Beth, now facing her fourth information set, keeps. Therefore Ann has a payoff of 1. With probability $\frac{4}{14}$, Ann gets a card of value 2 and Beth one of value 1. Ann keeps her card, and Beth faces her first information set, where she has to swap. The expected payoff for Ann is $\frac{1}{2}$. With probability $\frac{3}{14}$ both players get cards of value 2. Then neither one swaps and Ann's payoff is 0. Overall, the expected value for Ann when Ann plays strategy SK and Beth plays strategy SKSK is

$$\frac{3}{14} \cdot \frac{-2}{3} + \frac{4}{14} \cdot 1 + \frac{4}{14} \cdot \frac{1}{2} + \frac{3}{14} \cdot 0 = \frac{4}{14} = \frac{2}{7}.$$

Performing such calculations for every pair of pure strategies, we obtain the matrix of Ann's payoffs

	KKKK	KKKS	KSKS	KKSS	SKKK	SKKS	SKSK	SKSS
KK	0	0	0	0	$\frac{-2}{7}$	$\frac{-2}{7}$	$\frac{-2}{7}$	$\frac{-2}{7}$
SK	$\frac{4}{7}$	$\frac{3}{7}$	$\frac{3}{7}$	$\frac{2}{7}$	$\frac{3}{7}$	$\frac{2}{7}$	$\frac{2}{7}$	$\frac{1}{7}$

For the smaller extensive form in Figure 22.12, we get a 2×2 bimatrix.

	K	S
K	0	$\frac{-2}{7}$
S	$\frac{2}{7}$	$\frac{1}{7}$

Since the backward induction analysis of the subgames indicates that Beth should swap in her third and fourth information sets (after Ann has swapped), Beth's pure strategies in the smaller extensive form correspond to KKSS and SKSS in the larger form.

To consider larger examples, let us extend the definition of Le Her:

LE HER*(S, r) is defined as LE HER*(2, 4) except that we play with r cards of value 1, r cards of value 2, and so on, up to r cards of value S . Swapping is not allowed with the highest value card.

Allowing swapping a highest-valued card would not change the game, since nobody would give it away. The restriction reduces complexity (if only slightly). The original Le Her game had another restriction: if a player decided to swap but would get a highest value card, then the swap is not executed. This classical Le

Her game was played with thirteen values and with four duplicates of each. Bewersdorff ([B2004], Chapter 39) discusses it.

Let us now concentrate on LE HER*(3, 4).
Again we can cut off some subtrees and analyze them separately. Figure 24.2 shows the pruned extensive form: white diamonds are terminal vertices and indicate the subgames.




Figure 24.2. LE HER*(3, 4)

In the extensive form (which already has seen considerable work) Ann has two information sets with two options—holding a card of value 1 or 2, and Beth also has two information sets with two options. Therefore both players have four pure strategies, KK, SK, KS, and SS. The normal form is

	KK	SK	KS	SS
KK	0	$-\frac{8}{33}$	0	$-\frac{8}{33}$
SK	$\frac{8}{33}$	$\frac{4}{55}$	$\frac{34}{165}$	$\frac{2}{55}$
KS	$-\frac{2}{55}$	$-\frac{8}{55}$	$-\frac{2}{55}$	$-\frac{8}{55}$
SS	$\frac{34}{165}$	$\frac{28}{165}$	$\frac{28}{165}$	$\frac{2}{15}$

In the compressed version of the extensive form of LE HER*(S, r), Ann has S − 1 information sets with two options in each. Beth also has S − 1 information sets, where Ann does not swap. Thus Ann and Beth both have 2^{S−1} pure strategies. So we get an 8 × 8 matrix for LE HER*(4, r), a 16 × 16 matrix for LE HER*(5, r), and a 4096 × 4096 matrix for LE HER*(13, r), even for compressed versions. At this juncture you probably don’t want to discuss LE HER*(13, r).

24.3 Using Tools from Simultaneous Games for the Normal Form

The tools we have applied to simultaneous games can be applied to games in normal form. We replace the word “move” with “strategy”. Each player simultaneously selects a pure strategy and submits it to the referee, who then determines outcomes and payoffs. Let us use some of our tools on LE HER*(2, 4) and LE HER*(3, 4).

We produced several slightly different extensive forms for LE HER*(2, 4), and got different normal forms. Let us discuss the larger one, approximating fractions by decimals.

	KKKK	KKKS	KKSK	KKSS	SKKK	SKKS	SKSK	SKSS
KK	0	0	0	0	-0.29	-0.29	-0.29	-0.29
SK	0.57	0.43	0.43	0.29	0.43	0.29	0.29	0.14

Ann’s maximin strategy SK guarantees a payoff of 0.14 for Ann, and Beth’s maximin strategy SKSS guarantees a payoff of -0.14 for Beth. Beth’s payoffs are the negatives of the numbers shown, since the game is zero-sum. Ann’s second strategy SK strictly dominates KK. As for Beth, her strategy SKSS strictly dominates the first four, KKKK, KKKS, KSKS, and KKSS, and weakly dominates SKKK, SKKS, and SKSK. There is a Nash equilibrium, SK for Ann and SKSS for Beth, with an expected payoff of 0.14 for Ann. Ann swaps a value 1 card, and must, by the rules, keep a value 2 card. Beth swaps a card of value 1 and keeps a card of value 2 provided Ann kept her card. In both single-vertex information sets where Ann exchanged cards and Beth got a 1, Beth swaps (entries 3 and 4 in her strategy).

For LE HER*(3, 4) we have a simplified extensive form with rounded decimals:

	KK	SK	KS	SS
KK	0	-0.24	0	-0.24
SK	0.24	0.07	0.21	0.04
KS	-0.04	-0.15	-0.04	-0.15
SS	0.21	0.17	0.17	0.13

Ann’s maximin strategy is SS, with a guaranteed payoff of 0.13. Beth’s maximin strategy is also SS, with a guaranteed payoff of -0.13. This solves the game as the strategies form a Nash equilibrium. Ann swaps 1s and 2s. Beth does the same if Ann did not swap cards. If Ann swapped, Beth follows the reasonable path of-swapping if she got a 1 from Ann, and swapping if she got a 2 from Ann and Ann got a 3 from Beth. Ann has an advantage, with an expected payoff of 0.13.

24.4 Subgame Perfectness

In Chapter 22 we showed how subgames can be used to reduce the extensive form of a game. The resulting game has a normal form smaller than the normal form of the original extensive form. Since we are looking at different representations of the same game, we should expect the same solutions.

Let us look at the game with imperfect information shown in Figure 24.3. Ann moves first. In one case the game is over; in the other case Ann and Beth move simultaneously. Ann has two information sets with two moves in each. Her pure strategies are A_1A_3 , A_1A_4 , A_2A_3 , and A_2A_4 . Beth has one information set and the pure strategies B_1 and B_2 . The normal form is

	B_1	B_2
A_1A_3	0, 2	0, 2
A_1A_4	0, 2	0, 2
A_2A_3	3, 1	-2, -1
A_2A_4	1, -2	-3, -1




Figure 24.3. An example game




Figure 24.4. A subgame

There are three pure Nash equilibria: $A_1 A_3$ versus B_2 , $A_1 A_4$ versus B_2 , and $A_2 A_3$ versus B_1 . The extensive form has a subgame shown in Figure 24.4. This is a simultaneous game with the bimatrix:

	B_1	B_2
A_3	3, 1	-2, -1
A_4	1, -2	-3, -1

It has one Nash equilibrium, A_3 versus B_1 with payoffs 3 and 1.

Therefore we could cut off the subgame in the original extensive form, transforming the cut vertex (indicated by the white diamond) into a terminal vertex with payoff 3 and 1, as shown in the tree trunk in Figure 24.5. Ann will choose A_2 , with payoffs of 3 for Ann and 1 for Beth. This differs from the normal form calculation above, where we had three possible outcomes. The Nash equilibria $A_1 A_3$ versus B_2 and $A_1 A_4$ versus B_2 will not occur. Now we have a problem.




Figure 24.5. The trunk game

This problem was resolved by Reinhard Selten, who described how and why, when looking at subgames, we consider only some of the Nash equilibria. $A_1 A_3$ versus B_2 is a Nash equilibrium since Ann’s strategy $A_1 A_3$ is the best response to Beth’s strategy B_2 and vice versa. But a pre-game agreement on this would not be self-enforcing. When Beth decides, she already knows whether Ann has played A_1 or A_2 . If Ann deviates and starts with A_2 , Beth will also deviate and play B_1 . Ann can gain by deviating, since she can force Beth to deviate too.

A Nash equilibrium of the normal form is **subgame perfect** if it is a Nash equilibrium for every subgame. In our example, if we apply $A_1 A_3$ versus B_2 to the subgame in Figure 24.4 we get A_3 versus B_2 , which is not a Nash equilibrium for the subgame. The same holds for $A_1 A_4$ versus B_2 . Only $A_2 A_3$ versus B_1 , since it results in the Nash equilibrium A_3 versus B_1 in the subgame, is subgame perfect. It turns out that by using the subgame cutting-off method described in Chapter 22, we get exactly the subgame perfect Nash equilibria.

Historical Remark The concept of subgame-perfect Nash equilibria was introduced by the German economist Reinhard Selten in 1965. For this and other more complex achievements in game theory, Selten received the Nobel prize (shared with John Harsanyi and John Nash) in 1994. Selten was also one of the first to conduct experiments in game theory.

Reinhard Selten was born 1930 in Breslau, now Wroclaw. Growing up half Jewish in Hitler’s Germany was difficult and made Selten interested in politics and economics. After the war, he studied mathematics in Frankfurt, obtained a Ph. D in 1961, and became interested in the new topic of game theory. Among his important contributions was another refinement of a Nash equilibrium that he introduced in 1975: trembling hand perfectness. He was professor in Berlin, Bielefeld, and Bonn.

24.5 Special Case of Sequential Games with Perfect Information




Figure 24.6. A sequential game

Since we have the powerful tool of backward induction in sequential games with perfect information, there is no need to calculate the normal form. But they have a normal form. The strategies found by backward induction are exactly the subgame perfect Nash equilibria of the normal form. There may be more Nash equilibria (although they may not be very interesting). Take as an example the following game representing the situation in Europe in the 60s (as it was perceived by an unnamed military expert). In this game, the only solution found by backward induction is that the Soviets will attack, and the US will defend Europe conventionally. The Soviets at the time had two strategies, and the US also had two strategies—when attacked, responding with nuclear or with conventional weapons. The normal form looks like:

	nuclear	conventional
status quo	2, 2	2, 2
attack	0, 0	3, 1

It is easy to see that there are two Nash equilibria in pure strategies, namely (status quo, nuclear) with payoffs of (2, 2) and (attack, conventional) with payoffs of (3, 1). Only the second is the Zermelo-Kuhn solution found by backward induction. Only the second is subgame-perfect. This can be explained in terms of unconvincing threats: The US threaten nuclear retaliation if the Soviets attack, but the threat is not convincing since at the time the US must decide on a response (after the Soviets have chosen to attack), the US are better off choosing a conventional response.

Exercises

- Both Ann and Beth put \$1 in the pot. Ann gets a card from a stack of four queens and four kings and looks at it privately. Then Ann either folds, in which case Beth gets the money in the pot, or raises. Raising means that Ann has to put another dollar into the pot. When Ann has raised, Beth either folds, in which case Ann gets the pot, or Beth calls by putting one additional dollar into the pot. If Beth calls, Ann gets the pot if she has a king, otherwise Beth gets the pot.
 - Draw the extensive form of the game. How many pure strategies does Ann have, and how many pure strategies does Beth have?
 - Find the normal form of the game.
- Both Ann and Beth put \$1 into the pot. Ann gets a card and looks at it privately. Then Ann either checks, in which case Ann gets the money in the pot if Ann’s card is red, or Beth gets it if Ann’s card is black. Ann can also raise by putting another dollar into the pot. Now Beth either folds, in which case Ann gets the pot, or Beth calls by putting one additional dollar into the pot. If Beth calls, Ann gets the pot if she has a red card, otherwise Beth gets the pot.
 - Draw the extensive form of the game. How many pure strategies does Ann have, and how many pure strategies does Beth have?
 - Find the normal form of the game.

3. Ann starts a game by selecting (Ann doesn't draw, she chooses) two cards from a deck of cards containing four queens and four kings. Ann puts them face down in front of Beth. Beth is allowed to see one of them. Then Beth must guess whether the cards are two kings, two queens, or one king and one queen. If she is right, she wins \$1 from Ann, otherwise she has to pay \$1 to Ann.
 - Draw the extensive form of the game. How many pure strategies does Ann have, and how many pure strategies does Beth have?
 - Find the normal form of the game.
 - Eliminate weakly dominated strategies. Find the IEWD matrix, obtained by iterated elimination of weakly dominated strategies.
4. Ann starts a game by selecting (Ann doesn't draw, she chooses) one card from a deck of cards containing four queens and four kings. Ann puts it face down in front of Beth. Beth is not allowed to see it, but is allowed to see one card chosen randomly from the remaining deck of seven cards. Then Beth must guess whether the card face down is a king or a queen. If she is right, she wins \$1 from Ann, otherwise she has to pay Ann \$1.
 - Draw the extensive form of the game. How many pure strategies does Ann have, and how many pure strategies does Beth have?
 - Find the normal form of the game.
 - Are there any weakly dominated strategies?
5. Ann starts a game by selecting (Ann doesn't draw, she chooses) two cards from a deck of cards containing four queens and four kings. Ann puts them face down in front of Beth. Beth is not allowed to see them, but is allowed to see one card chosen randomly from the remaining deck of six cards. Then Beth must guess whether the two cards are two kings, two queens, or a king and a queen. If she is right, she wins \$1 from Ann, otherwise she has to pay Ann \$1.
 - Draw the extensive form of the game. How many pure strategies does Ann have, and how many pure strategies does Beth have?
 - Find the normal form of the game.
6. The game whose extensive form is in Figure 24.7 is called STRIPPED-DOWN POKER. It is a card game, played with a deck of four kings and four queens. Describe the rules of the game, and find the normal form.




Figure 24.7. STRIPPED-DOWN POKER

7. Find the normal form of the variant of STRIPPED-DOWN POKER that is played with four queens and only three kings. Can you find the normal form for decks with q queens and k kings?
8. The game whose extensive form is in Figure 24.8 is called MYERSON POKER. It is a card game, played with a deck of four kings and four queens. Describe the rules of the game, and find the normal form.




Figure 24.8. MYERSON POKER

9. Find the normal form of the variant of MYERSON POKER that is played with with four queens and only three kings. Can you find the normal form for decks with q queens and k kings?
10. How many pure strategies does Ann have in the game whose extensive form is in Figure 24.9? How many pure strategies does Beth have? Describe one pure strategy for Ann and one for Beth. How many reduced pure strategies do Ann and Beth have?




Figure 24.9. Another sequential game with imperfect information

CHAPTER 25

Example: VNM POKER and KUHN POKER

Prerequisites: Chapters 2, 8, 12, 16, 22, and 24.

Finally, a chapter on poker! Now we will learn how to play it well and make a lot of money, right? If this was your thought, I have to disappoint you. Advice on poker is possible, but the game is again too complex for a mathematical analysis. What we will do instead is study several small variants that we can analyze. This chapter will introduce the games, and do some analysis. We will complete the analysis in Chapters 31, 36, and 37.

25.1 Description

The two families of games we describe and partially analyze in this chapter are classics. What I call VNM POKER was introduced by von Neumann and Morgenstern in their monograph [VNM1944]. KUHN POKER (in the variant (3, 1, 1, 2)) was introduced by Kuhn in [K1950].

Both games have four parameters, S, r, n , and m , with $m < n$. There are cards of value from 1 to S in a deck, and each value occurs r times. So there are $S \cdot r$ cards. There are two players, Ann and Beth. Each player randomly gets a card from the deck, looks at it, but doesn't show it to her opponent. The ante is m , meaning that both player put m dollars into the pot at the start of the game. Each player's bet can raise to n if the player puts $n - m$ additional dollars in.

VNMPOKER(S, r, m, n) Ann moves first by checking (playing for m) or raising (playing for n).

- If Ann checks, both cards are revealed and the player with the higher card gets the pot of $2m$, and thus has m more than at the start. In the case of a draw the money is split equally and each one has the same as at the start.
- If Ann raises, she increases her total bet to n . Then Beth has two options, folding or calling.
 - If Beth folds, Ann gets the pot money of $n + m$, i.e., wins m . Beth's card is not revealed.
 - If Beth calls, she increases her stake to n . Both cards are then revealed, and the player with the higher card gets the $2n$ in the pot; i.e., wins n . In case of a draw the money is split equally.

KUHNPOKER(S, r, m, n) These games extend VNMPOKER. If Ann checks, the players enter VNMPOKER with Ann and Beth playing reversed roles. Ann moves first by either checking or raising.

- If Ann checks, then Beth can check or raise.
 - If Beth checks, both cards are revealed and the player with the higher card wins the pot, splitting it equally in case of a draw.
 - If Beth raises, she increases the bet to n . Then Ann has two options, folding or calling.
 - * If Ann folds, Beth gets the pot of $n + m$. Ann's card is not revealed.
 - * If Ann calls, she increases her bet to n . Then both cards are revealed, and the player with the higher card gets $2n$, splitting the money in case of a draw.
- If Ann raises, the game proceeds as in VNMPOKER when Ann raises.

Both games are zero-sum. Before we find their reduced normal forms, let us show their extensive forms, since this may clarify the descriptions. We look first at the modules in Figures 25.1 and 25.2, which describe Ann's and Beth's moves. We combine modules to give the games' full extensive forms.




Figure 25.1. The VNM POKER module




Figure 25.2. The KUHN POKER module

VNMPOKER and KUHNPOKER start with the random move of giving cards to both players. Ann can get any of the cards $1, 2, \dots, S$, and Beth as well (provided $r > 1$), so there are S^2 different alternatives if $r > 1$, and $S \cdot (S - 1)$ alternatives if $r = 1$. The parameters S and r influence the shape of the game tree, the number of modules, and the probabilities for the options of the random move. The payoffs obviously depend on m and n .

The alternatives for the initial random move are not equally likely. The cases where Ann and Beth have cards of equal value have probability

$$p_{XX} = \frac{1}{S} \cdot \frac{r-1}{rS-1} = \frac{r-1}{S(rS-1)}.$$

Why? The probability for, say a king, is $\frac{r}{rS} = \frac{1}{S}$. Once a king has been removed from the deck, there are $r-1$ kings left in $rS-1$ cards, so the probability of drawing another king is $\frac{r-1}{rS-1}$. The probability that both events happen is the product of the probabilities. Similarly, the case where Ann and Beth draw cards of different values has the higher probability

$$p_{XY} = \frac{1}{S} \cdot \frac{r}{rS-1} = \frac{r}{S(rS-1)}.$$




Figure 25.3. VNM POKER(3, 1, 2, 3)

Figure 25.3 shows VNM POKER(3, 1, 2, 3) and Figure 25.4 shows KUHN POKER(2, 4, 2, 3). Can you identify the modules? The information sets span different modules and, like threads in a fabric, hold the whole game together. There are no subgames.

25.2 VNM POKER

Student Activity Play VNM POKER(13, 4, 2, 3), using an ordinary poker deck, in the [VNMPokerseq13](#) applet. Choose two of the computer opponents and play at least 20 rounds against each. Which is easier to beat?

In VNM Poker what should Ann’s and Beth’s strategies be? If $r > 1$ there are $S \cdot S$ different ways of dealing one card value to Ann and one to Beth. Ann can see only her card and has therefore S information sets. In each she can check or raise. Therefore Ann has 2^S pure strategies, all reduced, since she has only one move. We display the strategies by sequences of Rs and Cs— for instance RCCR for $S = 4$ means that Ann raises with the lowest and the highest card and checks with the two middle-valued cards.

What pure strategies does Beth have? If Ann checks, Beth doesn’t have a choice, so it suffices to consider the other case, when Ann raises. Since Beth sees her own card but not Ann’s, Beth has S information sets. In each she will either call or fold. So she has 2^S pure strategies, all reduced. The strategies are encoded by sequences of Cs and Fs—FFCC for $S = 4$ means that she folds with the two lower-valued cards and calls with the two higher-valued cards.

To find Ann’s payoff if both players use a given pure strategy, we use the weighted average of Ann’s payoffs for the $S \cdot S$ card distributions, where the weights are the probabilities p_{XX} or p_{XY} . The payoffs are m , 0, or $-m$ in the case where Ann checks, m in case Ann raises and Beth folds, and n , 0, or $-n$ in the case where Ann raises and Beth calls. For example, if $S = 2$, Ann plays CR, and Beth CC, then Ann’s expected




Figure 25.4. KUHN POKER(2, 4, 2, 3)

payoff is

$$p_{XX} \cdot 0 + p_{XY} \cdot (-m) + p_{XY} \cdot n + p_{XX} \cdot 0,$$

with p_{XX} and p_{XY} depending on r and S . Beth's payoff is the negative of this, since the game is zero-sum.

To give another example, if $S = 3$, Ann plays RCR, and Beth plays FCC, then Ann's expected payoff is

$$p_{XX} \cdot m + p_{XY} \cdot (-n) + p_{XY} \cdot (-n) + p_{XY} \cdot m + p_{XX} \cdot 0 + p_{XY} \cdot (-m) + p_{XY} \cdot m + p_{XY} \cdot n + p_{XX} \cdot 0.$$

Fortunately some pure strategies are weakly dominated. When Ann holds a highest valued card, raising is not worse than checking. Any of Ann's strategies X...YC where Ann would check with a card of highest value S would be weakly dominated by the strategy X...YR that coincides with X...YC in all information sets except the one for the card value S , in which case Ann would raise. When Beth holds a highest valued card, calling dominates folding. Any of Beth's strategies X...YF with Beth folding with a card of value S would be weakly dominated by X...YC, the strategy identical to X...YF except for the information set holding a card of value S , in which case Ann would raise. Thus, when $S = 4$, out of the sixteen strategies for Ann and Beth, those that are not obviously weakly dominated are: for Ann, the eight strategies CCCR, CRRR, CRRC, CRRR, RCCR, RCRR, RRCR, and RRRR; and for Beth the eight strategies FFFC, FFCC, FCFC, FCCC, CFFC, CFCC, CCFC, and CCCC. The other strategies are weakly dominated.

Domination among Beth's strategies is even more frequent:

Theorem *All Beth's strategies except C...C and those of the form F...FC...C, starting with some Fs followed by some (at least one) Cs, are weakly dominated.*

Proof: We will show that a pure strategy for Beth (let's call it her original strategy, where she calls with card value i and folds with card value $i + 1$) is weakly dominated by the modified flipped strategy, where everything remains the same except that now Beth folds with a card value of i and calls with card value $i + 1$. For instance, CCFC is dominated by CFCC (the case $i = 2$), which itself is dominated by FCCC.

Out of the $S \cdot S$ terms for the card distributions that add to Beth's payoff, almost all are the same for the two strategies. The exceptions are those $2S$ summands corresponding to card values i and $i + 1$ for Beth. Let's call the term referring to a card of value x for Ann and a card of value y for Beth the (x, y) -term. Since the two card values i and $i + 1$ behave the same (both are larger or both are smaller) towards cards of value k different from i and $i + 1$, every (k, i) -term for one of the two strategies of Beth equals the corresponding $(k, i + 1)$ -term for the other strategy. Therefore all contributions to the payoff except the four where Ann and Beth both have cards of value i and $i + 1$ can be matched in both strategies. All that remains is to compare the four summands, the (i, i) -, $(i, i + 1)$ -, $(i + 1, i)$ -, and $(i + 1, i + 1)$ -term of Beth's payoff.

- If Ann checks for both i and $i + 1$, then the terms are $p_{XX} \cdot 0 + p_{XY} \cdot m + p_{XY} \cdot (-m) + p_{XX} \cdot 0$ independent of Beth's strategy.
- If Ann checks for i and raises for $i + 1$, then the terms are $p_{XX} \cdot 0 + p_{XY} \cdot m + p_{XY} \cdot (-n) + p_{XX} \cdot (-m) < 0$ in Beth's original strategy (where Beth calls with i and folds for $i + 1$), and $p_{XX} \cdot 0 + p_{XY} \cdot m + p_{XY} \cdot (-m) + p_{XX} \cdot 0 = 0$ in the flipped strategy (where Beth folds for i and calls for $i + 1$). The flipped strategy yields a higher payoff for Beth.
- If Ann raises for i and checks for $i + 1$, then the terms are $p_{XX} \cdot 0 + p_{XY} \cdot (-m) + p_{XY} \cdot (-m) + p_{XX} \cdot 0$ in Beth's original strategy, and $p_{XX} \cdot (-m) + p_{XY} \cdot n + p_{XY} \cdot (-m) + p_{XX} \cdot 0$ in the flipped strategy. The flipped strategy yields a higher payoff for Beth, since $p_{XX} \cdot (-m) + p_{XY} \cdot n > p_{XY} \cdot (-m)$.
- When Ann raises for both i and $i + 1$, the terms are $p_{XX} \cdot 0 + p_{XY} \cdot (-m) + p_{XY} \cdot (-n) + p_{XX} \cdot (-m)$ in Beth's original strategy, and $p_{XX} \cdot (-m) + p_{XY} \cdot n + p_{XY} \cdot (-m) + p_{XX} \cdot 0$ in the flipped strategy. The second expression is larger.

Thus for $S = 3$, only Ann's pure strategies CCR, CRR, RCR, and RRR are not necessarily dominated, and only Beth's three pure strategies FFC, FCC, and CCCC are not necessarily dominated. The matrix shows Ann's payoff for $S = 3, r = 4, m = 1, n = 4$:

	FFC	FCC	CCC
CCR	0.000	0.364	0.727
CRR	-0.273	0.000	0.727
RCR	-0.030	-0.273	0.000
RRR	-0.303	-0.636	0.000

Table 25.1. VNM POKER(3, 4, 1, 4), with some weakly dominated strategies already eliminated

The game has a Nash equilibrium of CCR versus FFC. Another is CCC versus FFC—though CCC is weakly dominated by CCR, but that doesn't prevent it from being part of a Nash equilibrium.

The most entertaining part of poker is bluffing, where a player raises though he or she has a low hand. Why does Ann not bluff? The reason seems to be that it is too risky. The raised bet of 4 is just too high

compared to the initial bet of 1. For smaller ratios n/m , we do not get a pure Nash equilibrium, and bluffing will occur there.

Now consider a second example with $S = 4$. Only the four pure strategies FFFC, FFCC, FCCC, and CCCC remain for Beth. The matrix, for $r = 4, m = 3$, and $n = 5$, is

	FFFC	FFCC	FFCC	CCCC
CCCR	0.000	$2/15 \approx 0.133$	$4/15 \approx 0.267$	$2/5 = 0.400$
CCRR	$1/60 \approx 0.017$	0.000	$4/15 \approx 0.267$	$8/15 \approx 0.533$
CRCR	$5/12 \approx 0.417$	$1/60 \approx 0.017$	0.000	$4/15 \approx 0.267$
CRRR	$13/30 \approx 0.433$	$-7/60 \approx -0.117$	0.000	$2/5 = 0.400$
RCCR	$49/60 \approx 0.817$	$5/12 \approx 0.417$	$1/60 \approx 0.017$	0.000
RCRR	$5/6 \approx 0.833$	$17/60 \approx 0.283$	$1/60 \approx 0.017$	$2/15 \approx 0.133$
RRCR	$37/30 \approx 1.233$	$3/10 = 0.300$	$-1/4 = -0.250$	$-2/15 \approx -0.133$
RRRR	$5/4 = 1.250$	$1/6 \approx 0.167$	$-1/4 = -0.250$	0.000

Table 25.2. VNM POKER(4, 4, 3, 5), with some weakly dominated strategies already eliminated

There is no domination among the strategies, and there is no pure Nash equilibrium. The maximin strategy for Ann is RCRR, with payoff guarantee of $1/60$, about 0.167, and the maximin strategy for Beth is FCCC, with a payoff guarantee of $-4/15$, about -0.267 (meaning a payoff of 0.267 for Ann). In the maximin strategy, Ann raises holding a card of 1, 3, or 4 but checks with a card of 2. Raising with a value of 1 can certainly be considered bluffing.

Excel sheets [VNMPoker2.xlsx](#), [VNMPoker3.xlsx](#), and [VNMPoker4.xlsx](#) allow you to create bimatrices like Table 25.2 for $S = 2, 3, 4$ and all parameters n, m , and r .

25.3 KUHN POKER

Student Activity Play KUHN POKER(3, 4, 2, 3) in the applet [KuhnPoker3](#). Choose two computer opponents and play at least 20 rounds against each. Which one is easier to beat?

Ann has $2S$ information sets in this game: Either she makes her first move, raise or check, or Ann has checked in her first move, Beth has raised, and Ann calls or folds. In each case she knows only the value of her card (between 1 and S). Beth has $2S$ information sets, determined by the value of her card (between 1 and S) and whether Ann has raised (the first S information sets) or checked (the second S information sets).

Therefore Ann has 2^{2S} pure strategies, indicated by S letters R or C for her choice of raising or checking, given her card from low to high, and then S letters F or C for folding or calling when she has checked at the beginning and Beth raised, again given Ann's card from low to high. For instance, for $S = 2$ we get the sixteen pure strategies CCFF, CCFC, CCCF, CCCC, CRFF, CRFC, CRCF, CRCC, RCFF, RCFC, RCCF, RCCC, RRCF, RRFC, RRCC, and RRCC. If Ann raises holding a given card, it doesn't matter if she would later check or fold with the same card. For this reason, the reduced strategy with an R in position x ($1 \leq x \leq S$) would have "•" in position $S + x$. In the case $S = 2$ the nine reduced pure strategies are CCFF, CCFC, CCCF, CCCC, CRF•, CRC•, RC•F, RC•C, and RR••.

Beth also has 2^{2S} pure strategies since she has $2S$ information sets— S where Ann has raised, and S where she has checked—with two options in each. For $S = 2$ they are FFCC, FFRC, FFRR, FCCC, FCCR, FCRC, FCRR, CFCC, CFRC, CFRR, CCCC, CCCR, CCRC, and CCRR. All are already reduced.

We can eliminate some of the (reduced) pure strategies because they are weakly dominated: If Ann holds a highest-value card, and checks while Beth raises, it is always better for Ann to call. If Ann were to fold,

she would lose m units, whereas she cannot lose when she calls. Therefore, for $S = 2$, Ann’s strategies CCFF, CCCF, and RC●F are weakly dominated. If Ann holds a highest-value card, raising does not dominate checking. It depends on Beth’s pure strategy. Assume Beth always folds when Ann raises and always raises when Ann checks. Then, when Ann holds a highest valued card but Beth does not, raising gives Ann a payoff of m units, but checking would give her a higher payoff of n .

If Beth holds a highest-value card, calling weakly dominates folding, and raising weakly dominates checking.

Let’s look into two examples for $S = 2$: KUHN POKER(2, 4, 2, 4) has a Nash equilibrium of CCFC versus FCCR. Compared with an initial bet of 2, an increased bet of 4 seems too high to try bluffing: Ann and Beth both check with a lower-value card. Beth always folds with a lower-value card, expecting Ann not to bluff. KUHN POKER(2, 4, 3, 4) has the more aggressive Nash equilibrium of RR●● versus CCRR.

	FCCR	FCRR	CCCR	CCRR
CCFC	0.00	0.14	0.00	0.14
CCCC	−0.57	0.00	−0.57	0.00
CRF●	0.00	−0.43	0.57	0.14
CRC●	−0.57	−0.57	0.00	0.00
RC●C	−0.14	0.43	−0.57	0.00
RR●●	−0.14	−0.14	0.00	0.00

Table 25.3. KUHN POKER(2, 4, 2, 4), with some weakly dominated strategies eliminated

	FCCR	FCRR	CCCR	CCRR
CCFC	0.00	−0.36	0.00	−0.36
CCCC	−0.29	0.00	−0.29	0.00
CRF●	0.00	−0.64	0.29	−0.36
CRC●	−0.29	−0.29	0.00	0.00
RC*C	0.36	0.64	−0.29	0.00
RR●●	0.36	0.36	0.00	0.00

Table 25.4. KUHN POKER(2, 4, 3, 4), with some weakly dominated strategies eliminated

For $S = 3$, we initially get a 18×16 matrix. However, six of Ann’s strategies and four of Beth’s are weakly dominated, and after eliminating them we arrive at the matrix in Table 25.5’

Exercises

1. Play some games against one of the computer players in [VNMPokerseq13](#) and describe the game of the computer player as you perceive it.
2. Use the applet [VNMPokerseq13CC](#) to play a tournament between the four computer players. Each should play 100 rounds against each other one. Which is the best player, based on this?
3. Find the maximin strategies and the security levels for KUHN POKER(3, 4, 2, 3), where part of the normal form is given in Figure 25.5.

	FFC CCR	FFC CRR	FFC RCR	FFC RRR	FCC CCR	FCC CRR	FCC RCR	FCC RRR	CCC CCR	CCC CRR	CCC RCR	CCC RRR
CCCFFC	0.00	−0.06	−0.55	−0.61	0.00	−0.06	−0.55	−0.61	0.00	−0.06	−0.55	−0.61
CCCFCF	−0.12	0.00	−0.06	0.06	−0.12	0.00	−0.06	0.06	−0.12	0.00	−0.06	0.06
CCRFF●	0.00	−0.18	−0.67	−0.85	0.12	−0.06	−0.55	−0.73	0.24	0.06	−0.42	−0.61
CCRFC●	−0.12	−0.12	−0.18	−0.18	0.00	0.00	−0.06	−0.06	0.12	0.12	0.06	0.06
CRCF●C	0.06	0.18	0.00	0.12	−0.12	0.00	−0.18	−0.06	0.00	0.12	−0.06	0.06
CRRF●●	0.06	0.06	−0.12	−0.12	0.00	0.00	−0.18	−0.18	0.24	0.24	0.06	0.06
RCC●FC	0.55	0.48	0.18	0.12	−0.06	−0.12	−0.42	−0.48	−0.24	−0.30	−0.61	−0.67
RCC●CC	0.42	0.55	0.67	0.79	−0.18	−0.06	0.06	0.18	−0.36	−0.24	−0.12	0.00
RCR●F●	0.55	0.36	0.06	−0.12	0.06	−0.12	−0.42	−0.61	0.00	−0.18	−0.48	−0.67
RCR●C●	0.42	0.42	0.55	0.55	−0.06	−0.06	0.06	0.06	−0.12	−0.12	0.00	0.00
RRC●●C	0.61	0.73	0.73	0.85	−0.18	−0.06	−0.06	0.06	−0.24	−0.12	−0.12	0.00
RRR●●●	0.61	0.61	0.61	0.61	−0.06	−0.06	−0.06	−0.06	0.00	0.00	0.00	0.00

Table 25.5. KUHN POKER(3, 4, 2, 3), with some weakly dominated strategies eliminated.

CHAPTER 26

Example: Waiting for Mr. Perfect

Prerequisites: Chapters 12, 16, 22, and 24.

WAITING FOR MR. PERFECT is played by one, two, or more players. It has awards, and every player knows what types there are, and how many of each type are available. There is a fixed number of rounds, and there are more awards than rounds. At the start of a round, one of the awards is selected at random. Players simultaneously indicate if they are interested in the award. The award is then given randomly, with equal probability, to one of those who expressed interest. Players who have won an award are out of the game in future rounds.

There are several variants of this game, depending on its parameters.

- How many players are playing? More players means tougher competition.
- How many rounds are played? The smaller the number of rounds relative to the number of players, the tougher the game becomes. For example, if there are n players and n rounds, everybody can win an award of some type.
- What kinds of awards are available, and what is their distribution? Is the population of awards small? Or is it huge (meaning that the distribution effectively does not change through the different rounds)?

In this chapter we assume that there are two players, two rounds, and three types of awards. For both players one type carries a payoff of 1; a second, a payoff of 2; and a third, a payoff of 3. We assume that in each round the award with payoff i occurs with probability p_i for $i = 1, 2$, and 3. We will use the normal form of the game to find its pure Nash equilibria.

Student Activity Play the game at least 20 times using cards, three 1s, two 2s, and one 3. In each round select a card at random by taking it off the top of a well-shuffled deck. Don't forget to replace the card after each round, to keep the probabilities fixed, and to shuffle between the rounds and games. After you are done, each player should write a strategy of how to play this game. Then play 20 more games, this time keeping track of the results. You can use the applet [Waiting2b](#) instead of cards. The applet also explains the title of the chapter.

26.1 The Last Round

Wer jetzt kein Haus hat, baut sich keines mehr.
Wer jetzt allein ist, wird es lange bleiben,
wird wachen, lesen, lange Briefe schreiben
und wird in den Alleen hin und her
unruhig wandern, wenn die Blätter treiben

— Rainer Maria Rilke

Who has not built his house, will not start now.
 Who is now by himself will long be so,
 Be wakeful, read, write lengthy letters, go
 In vague disquiet pacing up and down
 Denuded lanes, with leaves adrift below.

— translation by Walter Arndt

Players have a choice between expressing interest or rejecting an award. But in the last round, if we assume the awards have positive value, the players still in the game don't have a choice. Any player who had not yet won an award would express interest in the last round. The expected value of the second round award is $p_1 + 2p_2 + 3p_3$. If both players are still in the game without an award, since the award is won with probability $1/2$ (the losing player receiving nothing), the expected payoff for each would be $(p_1 + 2p_2 + 3p_3)/2$. If one player has already won an award, the expected payoff for the player remaining in the game would be $p_1 + 2p_2 + 3p_3$.

You may ask why we bother using p_1 , p_2 , and p_3 here and complicating our formulas. Why don't we select values, like $p_1 = 1/2$, $p_2 = 1/3$, and $p_3 = 1/6$? Our approach has an advantage: If we are later interested in different probability values, like $p_1 = 0.2$, $p_2 = 0.3$, and $p_3 = 0.5$, then we don't have to repeat the analysis, we would just give values to the parameters in the formulas we have derived. This method, shifting from the concrete to the general, is frequent in mathematics, and one of the reasons for its success and power.

26.2 The Eight Pure Strategies

Having established that the players don't have a choice in the second round, let's now focus on the first round. In it, the players may face an award of value 1, 2, or 3, and have two options in each case. Therefore each player has $2^3 = 8$ pure strategies, which are encoded by the three-letter words NNN, ..., III. The first letter indicates how the player would play if the first award would be a 1, with N indicating no interest and I indicating interest, and the same for an award of value 2 in the second entry, and an award of value 3 in the last. NII would be the strategy of expressing interest only for awards of value of 2 or 3 in the first round.

26.3 Computing the Payoffs

Since each player has eight pure strategies, the normal form of the game would be an 8×8 bimatrix. How are the payoffs computed? Since we want to discuss the general case, let's assume that Ann plays strategy $A_1 A_2 A_3$, where the A_i are either N or I, and Beth plays strategy $B_1 B_2 B_3$. We look at the payoffs for both players in the three cases where the first-round award is 1, 2, or 3 separately. Assume the first-round award is k , either 1, 2, or 3.

- If both are interested, i.e., if $A_k = B_k = I$, then the players will get the award with probability $1/2$, and with the remaining $1/2$ probability will have to be satisfied with the award in the second round, whose expected value is $p_1 + 2p_2 + 3p_3$. The expected payoff for each player is $k/2 + (p_1 + 2p_2 + 3p_3)/2 = (k + p_1 + 2p_2 + 3p_3)/2$.
- If only one of them, say Ann, is interested, i.e., if $A_k = I$ and $B_k = N$, then Ann gets the first round award with payoff k and Beth gets the second-round award with expected value $p_1 + 2p_2 + 3p_3$.
- If neither player is interested in a first-round award of value k , then they compete for the second-round award, and have expected payoffs of $(p_1 + 2p_2 + 3p_3)/2$.

Having computed the payoffs for Ann and Beth for each of the three cases of a first-round award k of 1, 2, or 3, we now combine the formulas. What we do is compute another expected value, weighting the payoffs

for the three cases by their probabilities p_1 , p_2 , and p_3 . For example, if Ann chooses strategy INI and Beth chooses strategy NII, then the expected payoff for Ann is

$$p_1 \cdot 1 + p_2(p_1 + 2p_2 + 3p_3) + p_3(3 + (p_1 + 2p_2 + 3p_3)/2),$$

and the expected payoff for Beth is

$$p_1(p_1 + 2p_2 + 3p_3) + p_2 \cdot 2 + p_3(3 + (p_1 + 2p_2 + 3p_3)/2).$$

The formulas have been put into the Excel sheet [Waiting1.xlsx](#). p_1 , p_2 can be typed into the black cells (p_3 is determined by them), and the normal form will be computed automatically.

Let's choose $p_1 = 1/2$, $p_2 = 1/3$, and $p_3 = 1/6$. Then the 8×8 bimatrix normal form is:

	III	IIN	INI	INN	NII	NIN	NNI	NNN
III	1.67, 1.67	1.78, 1.56	1.72, 1.61	1.823, 1.5	1.5, 1.83	1.61, 1.72	1.56, 1.78	1.67, 1.67
IIN	1.56, 1.78	1.42, 1.42	1.61, 1.72	1.47, 1.36	1.39, 1.94	1.25, 1.58	1.44, 1.89	1.31, 1.53
INI	1.61, 1.72	1.72, 1.61	1.33, 1.33	1.44, 1.22	1.44, 1.89	1.56, 1.78	1.17, 1.5	1.28, 1.39
INN	1.5, 1.83	1.36, 1.47	1.22, 1.44	1.08, 1.08	1.33, 2	1.19, 1.64	1.06, 1.61	0.92, 1.25
NII	1.83, 1.5	1.94, 1.39	1.89, 1.44	2, 1.33	1.42, 1.42	1.53, 1.31	1.47, 1.36	1.58, 1.25
NIN	1.72, 1.61	1.58, 1.25	1.78, 1.56	1.64, 1.19	1.31, 1.53	1.17, 1.17	1.36, 1.47	1.22, 1.11
NNI	1.78, 1.56	1.89, 1.44	1.5, 1.17	1.61, 1.06	1.36, 1.47	1.47, 1.36	1.08, 1.08	1.19, 0.97
NNN	1.67, 1.67	1.53, 1.31	1.39, 1.28	1.25, 0.92	1.25, 1.58	1.11, 1.22	0.97, 1.19	0.83, 0.83

26.4 Domination

Some pure strategies look odd. For instance, in strategy NIN the player would express interest in the first round only for an award of value 2. Why would a player show interest for a 2 but not for a 3? One could argue that a player playing this strategy avoids the fight in case of a 3, leaving the award to the other player, and then takes whatever comes in the second round. However, we will see in this section that NIN is a strongly dominated strategy and should therefore not be played.

In what follows we compare strategies for Ann differing only in how they handle one value k of the first-round award, with an eye toward seeing whether one of them dominates the other. Strategies are easy to compare, since in the two cases of having a first-round award different from k , the payoffs for Ann are identical no matter what Beth plays. Therefore to compare the payoffs for Ann, we have only to consider Ann's payoff in case the first round award equals k , and look at the two cases of Beth being interested in a value of k , or of not being interested. If one of Ann's strategies is better than the other in both cases, then we have discovered strong domination.

For what values of k should Ann show interest in the first round? Assume that $(k + p_1 + 2p_2 + 3p_3)/2 > (p_1 + 2p_2 + 3p_3)$ for a certain value k . That means that Ann's payoff if both are interested in k is higher than the expected payoff when Ann waits for the second round. If Beth is interested in value k , then it is also preferable for Ann to be interested in that value k . The inequality is equivalent to $k > p_1 + 2p_2 + 3p_3$. If Beth is not interested in a value k , then for Ann interest in k is better than no interest if $k > (p_1 + 2p_2 + 3p_3)/2$. Therefore, for all values $k > p_1 + 2p_2 + 3p_3$ every pure strategy where a player shows no interest for value k is strongly dominated by the corresponding pure strategy where the player is interested in value k . Since $p_1 + p_2 + p_3 = 1$, this is equivalent to

$$k - 1 > p_2 + 2p_3.$$

The condition never holds for $k = 1$, but always for $k = 3$ (assuming $p_3 < 1$). For $k = 2$ it depends.

There are also cases where not being interested for a value k is always preferable to being interested, behavior for all other values being equal. If Beth is interested in k , then Ann not being interested in k is better than being interested provided $p_1 + 2p_2 + 3p_3 > (k + p_1 + 2p_2 + 3p_3)/2$, i.e., if $p_1 + 2p_2 + 3p_3 > k$. If Beth is not interested in k , then Ann being not interested in k is better than being interested provided $(p_1 + 2p_2 + 3p_3)/2 > k$. Overall, the domination occurs provided $p_1 + 2p_2 + 3p_3 > 2k$. Since $p_1 + p_2 + p_3 = 1$, this is equivalent to

$$p_2 + 2p_3 > 2k - 1.$$

This is not true for $k = 2$ or 3 , but may occur for $k = 1$.

We summarize:

Theorem

- a) *The pure strategy $A_1 A_2 I$ always strongly dominates the pure strategy $A_1 A_2 N$. Players should always express interest in a value of 3 in the first round.*
- b) *If $1 > p_2 + 2p_3$, then the pure strategy $A_1 I A_3$ strongly dominates the pure strategy $A_1 N A_3$. Then players should express interest in a value of 2 in the first round.*
- c) *If $p_2 + 2p_3 > 1$, then the pure strategy $N A_2 A_3$ strongly dominates the pure strategy $I A_2 A_3$. In that case players should not express interest in a value of 1 in the first round.*

Depending on whether $p_2 + 2p_3$ is less than 1 or is between 1 and 2 we get two cases, but there are in both cases six strongly dominated pure strategies. If $p_2 + 2p_3 = 1$, we get four strongly dominated strategies and some weak domination. Let us look into the three exhaustive cases.

26.5 The Reduced Normal Forms in the Three Cases

26.5.1 The Case $p_2 + 2p_3 < 1$

After eliminating the strongly dominated strategies, we have the reduced bimatrix:

	III	NII
III	$A_{1,1}, B_{1,1}$	$A_{1,2}, B_{1,2}$
NII	$A_{2,1}, B_{2,1}$	$A_{2,2}, B_{2,2}$

The two strategies agree on expressing interest in the first round in case of an award of value 2 or 3. Thus, all expected values in all combinations have a common term of

$$\begin{aligned} M &:= p_2(2 + p_1 + 2p_2 + 3p_3)/2 + p_3(3 + p_1 + 2p_2 + 3p_3)/2 \\ &= (2p_2(p_2 + 1) + 3p_3(p_3 + 1) + p_1 p_2 + p_1 p_3 + 5p_2 p_3)/2. \end{aligned}$$

The coefficients are

$$\begin{aligned} A_{1,1} &= B_{1,1} = M + p_1(1 + p_1 + 2p_2 + 3p_3)/2 \\ A_{2,2} &= B_{2,2} = M + p_1(p_1 + 2p_2 + 3p_3)/2 \\ A_{1,2} &= B_{2,1} = M + p_1 \\ A_{2,1} &= B_{1,2} = M + p_1(p_1 + 2p_2 + 3p_3). \end{aligned}$$

Since $p_2 + 2p_3 < 1$, and therefore $p_1 + 2p_2 + 3p_3 < 2$, the coefficients are ordered as

$$A_{2,2} = B_{2,2} < A_{1,2} = B_{2,1} < A_{1,1} = B_{1,1} < A_{2,1} = B_{1,2}.$$

Ann's best response to Beth's strategy III is NII, and Ann's best response to Beth's strategy NII is III, and vice versa, since the game is symmetric. In this case there are two Nash equilibria in pure strategies: III — take everything in the first round— versus the more picky NII of rejecting 1 in the first round, and also the other way NII versus III. Essentially the reduced form is a game of CHICKEN, where expressing interest in the value of 1 (III) would be the Dove strategy, whereas the Hawk would reject a value of 1 in the first round and play NII.

26.5.2 The Case $p_2 + 2p_3 > 1$

As in the previous case, we eliminate strongly dominated strategies, and get the IESD bimatrix

	NII	NNI
NII	$A_{1,1}, B_{1,1}$	$A_{1,2}, B_{1,2}$
NNI	$A_{2,1}, B_{2,1}$	$A_{2,2}, B_{2,2}$

The two strategies agree on not expressing interest in the first round for an award of value 1, and on expressing interest for a value of 3. All expected values in all combinations have a common term of

$$\begin{aligned} N &:= p_1(p_1 + 2p_2 + 3p_3)/2 + p_3(3 + p_1 + 2p_2 + 3p_3)/2 \\ &= (p_1^2 + 3p_3(p_3 + 1) + 2p_1p_2 + 4p_1p_3 + 2p_2p_3)/2. \end{aligned}$$

and the coefficients are

$$\begin{aligned} A_{1,1} &= B_{1,1} = N + p_2(2 + p_1 + 2p_2 + 3p_3)/2 \\ A_{2,2} &= B_{2,2} = N + p_2(p_1 + 2p_2 + 3p_3)/2 \\ A_{1,2} &= B_{2,1} = N + p_2 \cdot 2 \\ A_{2,1} &= B_{1,2} = N + p_2(p_1 + 2p_2 + 3p_3). \end{aligned}$$

Since $p_2 + 2p_3 > 1$, and therefore $p_1 + 2p_2 + 3p_3 > 2$, the ordering of the coefficients is

$$A_{2,2} = B_{2,2} < A_{1,2} = B_{2,1} < A_{1,1} = B_{1,1} < A_{2,1} = B_{1,2}.$$

We get again a CHICKEN game with two pure Nash equilibria: Dove (NNI) versus Hawk (NII) and also the other way. As in the previous subcase, the Hawk refuses to accept a certain level of awards—in this case of value 2—which the Dove has to accept in the Nash equilibrium.

26.5.3 The Case $p_2 + 2p_3 = 1$

We have $p_1 + 2p_2 + 3p_3 = 2$. The IESD process leaves the reduced bimatrix

	III	INI	NII	NNI
III	$5/2, 5/2$	$5/2, 5/2$	$\frac{5-p_1}{2}, \frac{5+p_1}{2}$	$\frac{5-p_1}{2}, \frac{5+p_1}{2}$
INI	$5/2, 5/2$	$5/2 - p_2, 5/2 - p_2$	$\frac{5-p_1}{2}, \frac{5+p_1}{2}$	$\frac{5-p_1}{2} - p_2, \frac{5+p_1}{2} - p_2$
NII	$\frac{5+p_1}{2}, \frac{5-p_1}{2}$	$\frac{5+p_1}{2}, \frac{5-p_1}{2}$	$\frac{5-p_1}{2}, \frac{5-p_1}{2}$	$\frac{5-p_1}{2}, \frac{5-p_1}{2}$
NNI	$\frac{5+p_1}{2}, \frac{5-p_1}{2}$	$\frac{5+p_1}{2} - p_2, \frac{5-p_1}{2} - p_2$	$\frac{5-p_1}{2}, \frac{5-p_1}{2}$	$\frac{5-p_1}{2} - p_2, \frac{5-p_1}{2} - p_2$

It has nine Nash equilibria in pure strategies: The six pairs (NII,III), (NII,INI), (III,NII), (INI,NII), (III,NNI), (NNI,III) are Pareto-dominating the remaining three (NII,NII), (NII,NNI), and (NNI,NII).

Student Activity In the student activity at the beginning of the chapter, $p_1 = 1/2$, $p_2 = 1/3$, $p_3 = 1/6$. Therefore $p_2 + 2p_3 < 1$, so we have the two Nash equilibria of NII versus III. Did anybody write down one of these strategies? If so, how did you score? Were you more successful than the others with a different strategy than NII or III?

Project 43

Analyze the two player, two rounds variant with awards of value 1, 2, or 3, starting with three awards of value 1, two awards of value 2, and one award of value 3. In the first round one of the awards is chosen randomly, and in the second round one of the remaining five awards is selected randomly.

Project 44

How would the game change if the number of rounds is not fixed at the beginning? Instead after every round, either the game ends (say with probability $1/5$) or there is another round (with probability $4/5$).

Project 45

Discuss the two players and two rounds version with awards C, D, and E, where C is worth 2 for Ann and 3 for Beth, D is worth 3 for Ann and 1 for Beth, and E is worth 1 for Ann and 2 for Beth.

Project 46

Analyze the Broken Heart variant, where every player can express interest only once. If interest was expressed by both, the loser (in this round) can not continue in the second round and gets a payoff of 0. Use either fixed probabilities or analyze the game for general probabilities.

Project 47

Analyze the variant where expressing interest costs M . Use fixed probabilities.

Project 48

Analyze the 3-players 2-round version of the game, for fixed probabilities or for general probabilities.

CHAPTER27

Theory 8: Mixed Strategies

Student Activity Play ROCK-SCISSORS-PAPER against copro-robot Tarzan in the CoproRobot1a applet until you have won three times more than Tarzan.

27.1 Mixed Strategies

We have seen that Nash equilibria are likely outcomes of games. What happens in games without a Nash equilibrium in pure strategies? Even simple games like ROCK-SCISSORS-PAPER do not have a Nash equilibrium. In games with pure Nash equilibria players want to communicate their strategies to the other player(s) before the game, but in ROCK-SCISSORS-PAPER it is crucial to leave your opponent in the dark. You want to surprise your opponent, and that may be best achieved by surprising yourself. This could be done by delegating part of the decision about your strategy to a random device. This is the idea of a **mixed strategy**.

An example for a mixed strategy in ROCK-SCISSORS-PAPER is to play rock, scissors, or paper with probabilities 50%, 25%, or 25%. Before the game is played, the player decides randomly, based on these probabilities, which strategy to play. We can view pure strategies as mixed strategies where we play one of the options with probability 100% and the others with probability 0%.

If we add the mixed strategy (50% rock, 25% scissors, 25% paper) as an option for Ann in ROCK-SCISSORS-PAPER, then the expected payoffs for Ann against Beth’s pure strategies rock, scissors, paper, are 0, 0.25, −0.25 respectively, see the table below.

	Rock	Scissors	Paper
Rock	0	1	−1
Scissors	−1	0	1
Paper	1	−1	0
50-25-25 mix	0	0.25	−0.25

If Ann plays the mixed strategy and Beth plays rock, then with 50% probability there will be a tie (rock versus rock), with 25% probability Beth will win (Beth’s rock against Ann’s scissors), and with 25% probability Ann will win (paper against Beth’s rock). Thus the expected payoff for Ann when playing the mixed strategy against Beth’s rock is $50\% \cdot 0 + 25\% \cdot (-1) + 25\% \cdot 1 = 0$. Thus the values in the fourth row of the table are expected values of the corresponding values in the other rows and same column, using the probabilities of the mix. For instance, the second value in the fourth row is $50\% \cdot 1 + 25\% \cdot 0 + 25\% \cdot (-1) = 0.25$, and the third $50\% \cdot (-1) + 25\% \cdot 1 + 25\% \cdot 0 = -0.25$. Though the mixed strategy doesn’t dominate any of the pure strategies, it may be attractive to a player aiming at the maximin strategy since it guarantees a payoff of −0.25 compared to −1 in the other cases.

Beth is also entitled to mixed strategies. We assume that she chooses a mix of 25% rock, 50% scissors, 25% paper. Adding this mix as another one of Beth’s options, we obtain the bimatrix:

	Rock	Scissors	Paper	25-50-25 Mix
Rock	0	1	-1	0.25
Scissors	-1	0	1	0
Paper	1	-1	0	-0.25
50-25-25 mix	0	0.25	-0.25	0.0625

The new values are computed as before, as expected values, using the payoffs of the same row, weighted by the probability of the mix. For instance, the last entry in the first row is computed as $25\% \cdot 0 + 50\% \cdot 1 + 25\% \cdot (-1) = 0.25$. Even the payoff for 50-25-25 mix against 25-50-25 mix is computed this way, using the fourth row values, as $25\% \cdot 0 + 50\% \cdot 0.25 + 25\% \cdot (-0.25) = 0.0625$. Not too surprisingly, Ann’s mix, with its emphasis on rock, beats Beth’s mix, which is heavy on scissors.

27.1.1 Best Response

Student Activity In the applet [CoproRobot1b](#) find a good strategy to use against Jane who plays the 50% rock, 25% scissors, 25% paper strategy. Play until you have won three times more than Jane. I know you can do it.

As can be seen in the matrix above, paper is better against Ann’s 50-25-25 mix than rock or scissors, since it gives an expected payoff of 0.25 for Beth. The opponent emphasizes rock, so you must emphasize the move beating rock, which is paper. We could also look for best responses among the mixed strategies, but if we mix strategies of different value we would dilute the best response.

This is true in general. Since the payoffs of a mixed strategy are weighted averages of the payoffs for pure strategies, the maximum payoff is always achieved at a pure strategy.

Fact Among the best responses to mixed strategies of the other players there is at least one pure strategy.

The mixed strategies that are best responses are exactly all combinations of the best response pure strategies. To formulate this, the notion of **support** may be useful—the support of a mixed strategy are all pure strategies that occur with nonzero probability.

Theorem (Indifference Theorem) A mixed strategy of Ann is a best response to mixed strategies of the other players if each pure strategy in its support is a best response to the given mixed strategies of the other players.

So do we even need these fancy mixed strategies, if we can always react with old-fashioned pure strategies? We discuss this next.

27.1.2 Brown’s Fictitious Play

Assume a game is repeated. How would we find what mixed strategy the other player is using? He or she may not tell us. You may want to observe closely what the other player is doing, and react accordingly, by always playing the best response pure strategy to what you believe to be the other player’s mixed strategy.

The other player may observe you as well, believing your strategy is mixed, and reacting accordingly. That means you will adapt your play according to the play of the other player, and the other player adapts his

or her play according to your play. Even though both players think they are always playing a pure strategy, it looks as if they would play mixed strategies, since their pure strategies change. In their histories, they play the different pure strategies with some relative frequencies, which are perceived as probabilities by the other player. The relative frequencies change over time, but maybe eventually less and less—they might converge. This idea has been formulated by G.W. Brown [B1951]. He proposed that for two-person zero-sum games such a process will eventually converge and that the resulting mixed strategies are both best responses to each other.

This can be simulated in the Excel sheet [Brown.xlsm](#) or in [Brown10.xlsm](#). In ROCK-SCISSORS-PAPER, the mixes of the histories are 1/3 rock, 1/3 scissors, and 1/3 paper. Try it and see how long it takes until all percentages are really close to 1/3.

If you think that mixed strategies in Nash equilibria are always equal mixes of the pure strategies, then look at the following example:

Example 1 CRASH ROCK-SCISSORS-PAPER This is a simultaneous zero-sum game played by two players. Each one has three moves, rock, scissors, and paper. Scissors wins against paper, paper wins against rock, but rock wins big against scissors, crashing it with a loud sound, and giving a double payoff to the winner. It is ROCK-SCISSORS-PAPER, but with this payoff matrix:

	Rock	Scissors	Paper
Rock	0	2	−1
Scissors	−2	0	1
Paper	1	−1	0

Try the matrix in the Excel sheet. Both players will eventually play the 25%-25%-50% mix. Although scissors seems to be the weakest move, it is played more often than the other two moves.

Example 2 To give another example, this time for a non-simultaneous game, let us look at the normal form of VNM POKER(2, 4, 2, 3) from Chapter 25. After eliminating weakly dominated strategies, we get the matrix:

	FC	CC
CR	0	2/7
RR	1/7	0

Running it through our Excel sheet, we get a mix of 1/3 CR and 2/3 RR for Ann, and 2/3 FC and 1/3 CC for Beth. The mixed strategies translate into Ann raising with probability 2/3 when holding a low-value card, and always raising when holding a high-value card. Beth would call with probability 1/3 when holding a low-value card, and always when holding a high-value card. The expected payoff for Ann when both play their mix is 2/21.

Brown’s fictitious play is important even if you play only once. In a sense, both players would simulate thousands of rounds in their heads, and arrive at the mixed strategies. That’s where the word “fictitious” comes from. Even though you play the game just once and nobody can observe the probabilities you have chosen for the pure strategies, it is important to have a mixed strategy and use it to select your move.

27.1.3 Mixed Maximin Strategy, Mixed Security Level, and Linear Programs

Up to now, the maximin strategy was supposed to be a pure strategy. From now on we also allow mixed strategies.

Each of Ann’s mixed strategies achieves the smallest outcome (for Ann) against a pure strategy of Beth. The reasoning is similar as before: Ann’s payoff for a mixed strategy of Beth is a weighted average of Ann’s payoffs for the pure strategies of Beth, and it cannot be less than the smallest value.

For every one of Ann’s infinitely many mixed strategies, we next create the row of the payoffs versus the finitely many pure strategies of Beth. The **mixed maximin strategy** is the mixed strategy with the highest lowest entry in the row—it is the mixed strategy that guarantees the largest expected payoff, which we may call the **mixed security level** for Ann. If the game is a two-player zero-sum game, this is also called the **value**.

Example 3 Take **ROCK-SCISSORS-PAPER2** as example, where whenever Ann plays scissors the payoff is doubled. The payoff matrix is

	BRock	BScissors	BPaper
ARock	0	1	−1
AScissors	−2	0	2
APaper	1	−1	0

We show a few rows. The formula for calculating the entries in a row is given in the last row.

	BRock	BScissors	BPaper
ARock	0	1	−1
AScissors	−2	0	2
APaper	1	−1	0
$\frac{1}{3}, \frac{1}{3}, \frac{1}{3}$ -mix	−1/3	0	1/3
$\frac{1}{4}, \frac{1}{4}, \frac{1}{2}$ -mix	0	−1/4	1/4
$\frac{2}{5}, \frac{1}{5}, \frac{2}{5}$ -mix	0	0	0
p_1, p_2, p_3 -mix	$p_1 \cdot 0 + p_2 \cdot (-2) + p_3 \cdot 1$	$p_1 \cdot 1 + p_2 \cdot 0 + p_3 \cdot (-1)$	$p_1 \cdot (-1) + p_2 \cdot 2 + p_3 \cdot 0$

The smallest entries in the first six rows are −1, −2, −1, −1/3, −1/4, and 0. So, among the six strategies, the one with the largest guarantee for Ann is the $\frac{2}{5}, \frac{1}{5}, \frac{2}{5}$ -mix. But keep in mind that there are infinitely many possible rows. Might there be rows with better guarantees?

Let’s now move to a general 3 × 3 bimatrix.

	Left	Middle	Right
Up	$A_{1,1}, B_{1,1}$	$A_{1,2}, B_{1,2}$	$A_{1,3}, B_{1,3}$
Middle	$A_{2,1}, B_{2,1}$	$A_{2,2}, B_{2,2}$	$A_{2,3}, B_{2,3}$
Down	$A_{3,1}, B_{3,1}$	$A_{3,2}, B_{3,2}$	$A_{3,3}, B_{3,3}$

Ann’s task is to maximize the minimum of the three values $p_1 \cdot A_{1,1} + p_2 \cdot A_{2,1} + p_3 \cdot A_{3,1}$, $p_1 \cdot A_{1,2} + p_2 \cdot A_{2,2} + p_3 \cdot A_{3,2}$, and $p_1 \cdot A_{1,3} + p_2 \cdot A_{2,3} + p_3 \cdot A_{3,3}$, by choosing her probabilities p_1 , p_2 , and p_3 (non-negative, with sum 1). We could reformulate this by saying that

1. Ann has to choose four values, v , p_1 , p_2 , and p_3 , satisfying
2.

• $p_1 \geq 0$

• $p_2 \geq 0$

• $p_3 \geq 0$

• $p_1 + p_2 + p_3 = 1$

• $v \geq p_1 \cdot A_{1,1} + p_2 \cdot A_{2,1} + p_3 \cdot A_{3,1}$

• $v \geq p_1 \cdot A_{1,2} + p_2 \cdot A_{2,2} + p_3 \cdot A_{3,2}$

• $v \geq p_1 \cdot A_{1,3} + p_2 \cdot A_{2,3} + p_3 \cdot A_{3,3}$
3. such that v is maximized.

This is a **linear program**.

Using methods from linear programming (which we will not cover), one finds that in Example 3, $p_1 = 2/5$, $p_2 = 1/5$, $p_3 = 2/5$ is indeed the only mixed maximin strategy with mixed security level of 0. For Beth, the mixed maximin strategy is $q_1 = 1/3$, $q_2 = 1/3$, $q_3 = 1/3$. See [HL2010] for details.

27.2 Mixed Nash Equilibria

A Nash equilibrium of mixed strategies is called a **mixed Nash equilibrium**.

27.2.1 Two-player Zero-sum Games

These games have two nice features:

Theorem (von Neumann (1928)) *Every finite two-person zero-sum game has at least one Nash equilibrium of mixed strategies. They are the maximin mixed strategies.*

This theorem follows from von Neumann’s famous Minimax Theorem, which proves among other things the existence of each player’s (mixed) maximin strategy in a finite two-player zero-sum game. Such a Nash equilibrium can also be obtained by Brown’s fictitious play process.

Theorem (Julia Robinson) [R1951] *If two players play a zero-sum game in normal form repeatedly, and if in each round a player chooses the best response pure strategy against the observed mixed strategy of the other player, then the mixed strategies converge to a pair of mixed strategies forming a Nash equilibrium.*

27.2.2 Non-Zero-sum Games

Things become more complicated. Look at the following two examples, one zero-sum, the other not:

Example 4

Game 1:	Left	Right	Game 2:	Left	Right
Up	0, 0	10, −10	Up	0, 0	10, 5
Down	5, −5	0, 0	Down	5, 10	0, 0

In corresponding situations Ann's payoffs are the same (but not Beth's).

Let's start with the zero-sum game 1: It can be shown that Ann's optimal mixed strategy is to choose up with probability $1/3$ (and therefore down with probability $2/3$), and Beth best chooses left with probability $2/3$. The expected payoff for Ann is $10/3$, and therefore $-10/3$ for Beth.

Therefore this strategy is the maximin strategy for Ann if mixed strategies are allowed. Since the maximin strategy and value depend only on Ann's payoffs, which don't change between game 1 and game 2, the strategy with $1/3$ up and $2/3$ down is also the maximin strategy for Ann in game 2 (with mixed strategies allowed). Since game 2 is symmetric, Beth has a similar maximin strategy of choosing $1/3$ left and $2/3$ right. The expected payoff for both when they play this way is $10/3$. This is not a Nash equilibrium in game 2—every deviation of Ann in the direction of more up is rewarded, as is every deviation of Beth in the direction of more left. We will show that, if mixed strategies are allowed, game 2 has three mixed Nash equilibria:

- In the first, Ann chooses up with probability $2/3$ and Beth chooses left with probability $2/3$, giving the same expected payoff of $10/3$ for both, but no reason to deviate.
- In the second Ann chooses up and Beth chooses right. The payoffs are 10 for Ann and 5 for Beth.
- The last Nash equilibrium is where Ann chooses down and Beth left. The payoffs are 5 and 10, respectively.

Although the maximin strategies do not necessarily form a mixed Nash equilibrium for general games, mixed Nash equilibria always exist, even for non-zero-sum games with an arbitrary number of players:

Theorem (John F. Nash 1950) *Every finite game has at least one Nash equilibrium in pure or mixed strategies.*

27.3 Computing Mixed Nash Equilibria

It is easy to check whether a set of mixed strategies, one for each player, forms a Nash equilibrium. According to the Indifference Theorem, all we have to do is the following for every player, say Ann: We need to check whether the pure strategies that occur in Ann's mix with nonzero probability are best responses to the mixed strategies of the other players. But the question now is how to find a Nash equilibrium. In general this is not easy, even when we have the normal form. Without the normal form, if the game is described by its extensive form, even calculating the normal form may be not feasible, since it may be far too large. See Chapter 34 for an example. In the following we assume that the normal form is given.

What we should do first is eliminate all strictly dominated pure strategies from the normal form.

Fact *A strictly dominated pure strategy never occurs in the support of a Nash equilibrium mixed strategy, but a weakly dominated strategy may.*

As mentioned, for two-player zero-sum games, the problem of finding Nash equilibria in mixed strategies can be formulated as a linear programming problem. Around 1947 George Dantzig developed a method for solving them, the Simplex Algorithm. It is in practice quite fast, although there are a few artificial cases where it takes a long time to find the solution. For example, there are linear programs with n variables whose solution requires 2^n iterations. So the running time of the simplex algorithm is exponential. [C1983], p. 124,

pp 47–49. Karmarkar and others have developed algorithms for solving linear programming problems that find solutions in polynomial time. This means that the solution time is bounded by a polynomial whose degree is based on the size of the problem. [HL2010].

For general games, if the numbers of rows and columns are small, the Indifference Theorem can be used to find Nash equilibria. Then solving a few systems of linear equations give all mixed Nash equilibria. We will do this for two-player games and bimatrices of sizes 2×2 and 3×3 .

This method becomes infeasible if the number of options for the players is large. In a two-player game where the normal form has 10 rows and 10 columns, we would need to solve systems with up to 10 equations and 10 variables, which is not too difficult using technology. But we would have to solve more than 1,000,000 of such systems, which would take a long time.

More sophisticated approaches are needed. For two-player games, the best approach found so far is an algorithm developed by Lemke and Howson (who by the way also showed that for zero-sum two-player games the number of Nash equilibria is odd except in degenerate cases). Still, like the simplex algorithm, it sometimes has exponential running time [LH1964].

Brown’s fictitious play can be used to get approximate values of mixed Nash equilibria. However, the method does not always work, and when it does work, convergence may be slow—even 50000 iterations may not be enough to be within 1% of the Nash equilibrium probabilities. Further, having found some Nash equilibrium using Brown’s fictitious play, how do we know that there are not more Nash equilibria?

So you see that games are not always easy to solve.
We will discuss methods for solving several small game types.

27.3.1 Small Two-player Zero-sum Games (optional)

2 × n zero-sum games

Let’s look at the example

Example 5

	B_1	B_2	B_3	B_4
A_1	1.25	0.5	2	3
A_2	2	3.5	1	0.5

When she plays this game, Ann chooses from two options, with the probability p for her first option and probability $1 - p$ for her second. Using the mixed strategy, against any of Beth’s pure strategies Ann will have a payoff that is a linear function in p , with $p = 0$ and $p = 1$ giving the payoffs in the pure cases, as illustrated in the next table.

	B_1	B_2	B_3	B_4
A_1	1.25	0.5	2	3
Mix with $p = 3/4$	0.875	2.375	1.75	2.375
Mix with $p = 1/2$	1.25	2.75	1.5	1.75
Mix with $p = 1/4$	1.625	3.125	1.25	1.125
A_2	2	3.5	1	0.5

For Beth’s i th move we draw the straight line from $(0, A_{1,i})$ to $(1, A_{2,i})$. We then draw a curve, starting on the left on the line closest to the x -axis. When lines cross, we jump and follow the line nearest to the x -axis. That is, we mark the function defined as the minimum of the four linear functions. Its largest value is the security level for Ann, and the corresponding p -value belongs to Ann’s Nash equilibrium mix.





Figure 27.1. Graph for Example 5, a 2×4 zero-sum game

The straight lines are drawn in Figure 27.1. The Nash equilibrium p is about 0.55, and Ann’s guaranteed payoff is a little more than 1.5. Since the point lies at the intersection of the straight lines belonging to B_1 and B_3 , Beth will mix these two pure strategies.

3 $\times$ n zero-sum games

A similar approach works for the $3 \times n$ case, but it requires three dimensions and is harder to compute and follow. Ann’s mixed strategy would involve probabilities p_1, p_2, p_3 with $p_1 + p_2 + p_3 = 1$, p_i being the probability that Ann plays her i th strategy. A triple of numbers corresponds to a point on a triangle, using barycentric coordinates. They are calculated as the ratio of shortest distance of the point to one side to the shortest distance of the third point of the triangle to the same side. Now we construct straight lines perpendicular to the triangle at the three points, and put a triangle on these pillars at heights $A_{1,i}, A_{2,i}, A_{3,i}$ for every pure strategy i of Beth. The triangles intersect. The point on the triangle having most space above it until it meets the lowest ceiling is Ann’s maximin mix, and the height to the lowest ceiling there is Ann’s security level for mixed strategies.

Example 6 In the example in Figure 27.2 the triangle and the three pillars are drawn in white. We are looking up, towards the three slanted ceilings, from below the white triangle.



	B_1	B_2	B_3
A_1	5	7	4
A_2	3	5	6
A_3	6	4	5

Figure 27.2. Graph for Example 6, a 3×3 zero-sum game

The matrix is that of CRASH ROCK SCISSORS PAPER with 5 added to Ann’s payoffs. The Nash equilibrium is located at the point where the three planes intersect.

27.3.2 Solving Small non Zero-sum Two-player Games by Solving Equations (optional)

The simplest normal forms are for two-player games where each player has exactly two strategies. Here is an example.

Example 7

	Left	Right
Up	1, 3	3, 2
Down	4, 1	2, 4

We say that Ann mixes if she plays a mixed strategy $(p, 1 - p)$ with $0 < p < 1$, and Beth mixes if she plays $(q, 1 - q)$ with $0 < q < 1$. There are three cases of mixed Nash equilibria: where Ann mixes and Beth doesn't, where Beth mixes and Ann doesn't, and where both mix.

Let us look at the third case first. Assume that Ann chooses up with probability p , $0 < p < 1$ and down with probability $1 - p$. In the same way, Beth chooses left with probability q , $0 < q < 1$ and right with probability $1 - q$. The Indifference Theorem implies that Ann has the same expected payoffs when playing up or down provided Beth keeps mixing. They are $q \cdot 1 + (1 - q) \cdot 3$ and $q \cdot 4 + (1 - q) \cdot 2$. Therefore

$$q \cdot 1 + (1 - q) \cdot 3 = q \cdot 4 + (1 - q) \cdot 2.$$

Similarly, in a Nash equilibrium Beth has the same payoffs with each of her options when Ann mixes. Therefore

$$p \cdot 3 + (1 - p) \cdot 1 = p \cdot 2 + (1 - p) \cdot 4.$$

We have two linear equations with two variables p and q . They are simpler than the general case, since p does not occur in the first equation and q not in the second. We can solve both equations separately to get $q = \frac{1}{4}$ and $p = \frac{3}{4}$.

What if one player mixes, and the other doesn't? Assume Ann mixes, plays up with probability p and down with probability $1 - p$, but that Beth plays the pure strategy left. Then the Indifference Theorem implies that both up and down are best responses to Beth's left, which is not the case, since Ann's payoffs, 1 and 4, are not equal. If we had a Nash equilibrium, then up versus left and down versus left would both have to be pure Nash equilibria. But since this example doesn't contain any pure Nash equilibria, we can exclude all cases where one player mixes and one doesn't.

The example has just one Nash equilibrium, $\frac{3}{4}$ up and $\frac{1}{4}$ down versus $\frac{1}{4}$ left and $\frac{3}{4}$ right.

The same method can be applied to larger examples, but the number of systems of equations we may have to solve increases. Take the example:

Example 8

	Left	Middle	Right
Up	3, 4	1, 2	4, 1
Middle	4, 2	2, 2	1, 4
Down	3, 1	4, 3	2, 4

We have one case where Ann mixes all three strategies, three cases where Ann mixes two of them, and three cases where she uses a pure strategy, and the same holds for Beth. If we pair Ann's seven cases with Beth's seven cases we get 49 possible patterns for Nash equilibria.

Let us start with the most interesting pattern, where both Ann and Beth mix three strategies. Assume that Ann chooses up with probability p_1 , middle with probability p_2 , and down with probability $1 -$

$p_1 - p_2$. Beth chooses left with probability q_1 , middle with probability q_2 , and right with probability $1 - q_1 - q_2$. The Indifference Theorem gives us two double equations, namely

$$\begin{aligned} 3q_1 + 1q_2 + 4(1 - q_1 - q_2) &= 4q_1 + 2q_2 + 1(1 - q_1 - q_2) = 3q_1 + 4q_2 + 2(1 - q_1 - q_2) \\ 4p_1 + 2p_2 + 1(1 - p_1 - p_2) &= 2p_1 + 2p_2 + 3(1 - p_1 - p_2) = 1p_1 + 4p_2 + 4(1 - p_1 - p_2). \end{aligned}$$

Each can be broken into two equations. So we have a system of four linear equations in the four variables p_1, p_2, q_1 , and q_2 :

$$\begin{aligned} 3q_1 + 1q_2 + 4(1 - q_1 - q_2) &= 4q_1 + 2q_2 + 1(1 - q_1 - q_2) \\ 3q_1 + 1q_2 + 4(1 - q_1 - q_2) &= 3q_1 + 4q_2 + 2(1 - q_1 - q_2) \\ 4p_1 + 2p_2 + 1(1 - p_1 - p_2) &= 2p_1 + 2p_2 + 3(1 - p_1 - p_2) \\ 4p_1 + 2p_2 + 1(1 - p_1 - p_2) &= 1p_1 + 4p_2 + 4(1 - p_1 - p_2). \end{aligned}$$

Collecting terms, we get

$$\begin{aligned} 2q_1 + 4q_2 &= 3 \\ 2q_1 + 5q_2 &= 2 \\ 4p_1 + 2p_2 &= 2 \\ 6p_1 + 1p_2 &= 3. \end{aligned}$$

The variables q_1 and q_2 occur in the first and the second equations, and p_1 and p_2 in the third and the fourth, so the first two equations can be solved separately from the last two. We get $q_1 = 3.5, q_2 = -1, p_1 = 0.5, p_2 = 0$, which is not a valid solution because p_1, p_2, q_1 , and q_2 are probabilities.

The next case is where Ann mixes two strategies, say up and middle with probabilities p and $1 - p$, and Beth mixes all three. We get three equations (one equation and a double equation) with three variables:

$$\begin{aligned} 3q_1 + 1q_2 + 4(1 - q_1 - q_2) &= 4q_1 + 2q_2 + 1(1 - q_1 - q_2) \\ 4p + 2(1 - p) &= 2p + 2(1 - p) = 1p + 4(1 - p). \end{aligned}$$

The first equation contains two variables, and the second and third (stemming from the double equation) contain only one. So we may not have a solution, but if we have one we have infinitely many. In our case, $4p + 2(1 - p) = 2p + 2(1 - p)$ implies $p = 0$, and $2p + 2(1 - p) = 1p + 4(1 - p)$ implies $p = \frac{2}{3}$, a contradiction, so there is no solution.

If we assume that Ann plays up with probability p and down with probability $1 - p$, and that Beth mixes all three options, we get

$$\begin{aligned} 3q_1 + 1q_2 + 4(1 - q_1 - q_2) &= 3q_1 + 4q_2 + 2(1 - q_1 - q_2) \\ 4p + 1(1 - p) &= 2p + 3(1 - p) = 1p + 4(1 - p). \end{aligned}$$

They give $p = \frac{1}{2}$ and triples $(q_1, \frac{2-2q_1}{5}, \frac{3-3q_1}{5})$, with $0 \leq q_1 \leq 1$, as solutions.

The other four cases where one player mixes two pure strategies and the other three lead to no solution, and therefore to no further Nash equilibrium.

In the same way, a Nash equilibrium where one player mixes and the other plays a pure strategy can occur only if we have two pure Nash equilibria in the same row, or two in the same column. Since our example doesn't have any pure Nash equilibria, we can exclude these cases.

If both Ann and Beth mix two strategies, essentially the formulas for the 2×2 case can be used. Take as an example the case where Ann mixes up and middle, using up with probability p , and that Beth mixes left and right, using left with probability q . Then we get

$$3 \cdot q + 4 \cdot (1 - q) = 4 \cdot q + 1 \cdot (1 - q),$$

$$4 \cdot p + 2 \cdot (1 - p) = 1 \cdot p + 4 \cdot (1 - p).$$

with the solution $p = \frac{2}{5}$ and $q = \frac{3}{4}$.

Exercises

1. Find mixed Nash equilibria of the zero-sum game

	B_1	B_2	B_3
A_1	2	1	-2
A_2	1	-1	1
A_3	0	2	-1

2. Find mixed Nash equilibria of the game with normal form

	B_1	B_2	B_3	
A_1	2, 1	1, 3	2, 1	2, 1
A_2	1, 3	3, 2	1, 1	1, 2
A_3	2, 1	1, 3	3, 2	2, 3

3. Find mixed Nash equilibria of the game with normal form

	B_1	B_2	B_3	
A_1	1, 1	2, 2	3, 3	4, 1
A_2	2, 3	4, 3	1, 1	3, 2
A_3	3, 4	4, 1	2, 4	1, 4
A_2	3, 2	2, 4	4, 2	1, 3

4. Find all mixed Nash equilibria of the game

	Left	Right
Up	1, 4	2, 1
Down	4, 1	2, 3

using the method described in Section 27.3.2.

5. Find all mixed Nash equilibria of the game

	Left	Right
Up	1, 2	2, 3
Down	2, 1	2, 3

using the method described in Section 27.3.2.

6. For the game given by the normal form

	B_1	B_2	B_3
A_1	3, 2	1, 1	2, 3
A_2	2, 3	3, 2	1, 1
A_3	1, 1	2, 3	3, 2

- a) Try to find mixed Nash equilibria using Brown's fictitious play. When you think the numbers are stable, run 3000 more iterations. What do you observe?
- b) Find as many as possible Nash equilibria using the method described in Section 27.3.2.

7. For the game given by the normal form

	B_1	B_2	B_3
A_1	3, 3	1, 2	3, 1
A_2	4, 2	2, 2	1, 4
A_3	3, 1	3, 3	2, 3

- a) Try to find mixed Nash equilibria using Brown’s fictitious play. When you think the numbers are stable, run 3000 more iterations. What do you observe?

b) Find as many as possible Nash equilibria using the method described in Section 27.3.2.
8. Finish the analysis of Exercise 1 of Chapter 24 by finding all mixed Nash equilibria.
9. Finish the analysis of Exercise 2 of Chapter 24 by finding all mixed Nash equilibria.
10. Finish the analysis of Exercise 3 of Chapter 24 by finding all mixed Nash equilibria.
11. Finish the analysis of Exercise 4 of Chapter 24 by finding all mixed Nash equilibria.
12. Finish the analysis of Exercise 5 of Chapter 24 by finding all mixed Nash equilibria.
13. Finish the analysis of Exercise 6 of Chapter 24 by finding all mixed Nash equilibria.
14. Finish the analysis of Exercise 7 of Chapter 24 by finding all mixed Nash equilibria.
15. Finish the analysis of Exercise 8 of Chapter 24 by finding all mixed Nash equilibria.
16. Finish the analysis of Exercise 9 of Chapter 24 by finding all mixed Nash equilibria.

The next two projects refer to **balanced 3-spinners**, which are devices with three numbers on them whose sum is zero. If one spins one, each number appears with equal probability. So 3-spinners are like 3-sided dice.

Project 49

Balanced 3-spinner duel with five options Two players each possess five balanced 3-spinners, namely $(-4, -1, 5)$, $(-2, -2, 4)$, $(-1, 0, 1)$, $(-3, 1, 2)$, and $(0, 0, 0)$. They simultaneously choose one. Then they spin it. The player whose spinner shows the higher number wins 1 unit, which the other player loses. For a tie the payoff is 0 for both. How would the players play?

Project 50

Balanced 3-spinner duel Two players simultaneously choose a balanced 3-spinner with integers. Then each player spins it. The player whose spinner shows the higher number wins 1 unit, which the other player loses. For a tie the payoff is 0 for both.

Both players have infinitely many options. Can you still discuss how they should play?

Project 51

COLONEL BLOTTO(4, 9, 9) Analyze this game described in Chapter 21. Use the BlottoC applet.

Project 52

Iterated COLONEL BLOTTO This game in rounds starts with COLONEL BLOTTO(4, 9, 9), discussed in Chapter 21. The loser loses one army and COLONEL BLOTTO(4, 8, 9) or COLONEL BLOTTO(4, 9, 8) is played. In each round, a game of the form COLONEL BLOTTO(4, x , 9) or COLONEL BLOTTO(4, 9, x) with $x \leq 9$ is played. If a player with nine armies wins, the other loses one army. If a player with less than nine armies wins, he wins an army. In case of a draw, the number of armies is unchanged. Then another COLONEL BLOTTO game is played. The player who first is reduced to five armies loses the game.

This game can be played in the three applets [BlottoIter](#) (you against another human), [BlottoIterC](#) (you against the computer), or [BlottoIterCR](#) (two optimal computer players play 20 games against each other).

Discuss the features of the game. What kind of game is it, what kind of tools can be used? Is there anything special about it? Try to analyze it, using the Excel sheet [Blotto.xlsx](#). Try to answer the question of how much the outcome of the whole game is already decided when the first round has been played and one player has lost. This can be done using the analysis, or by doing experiments with the applet [BlottoIterCR](#).

Project 53

Simultaneous QUATRO-UNO Analyze the game Simultaneous QUATRO-UNO discussed in Project 8 in Chapter 6.

Project 54

Are there any mixed Nash equilibria in the Waiting for Mr. Perfect game discussed in Chapter 26?

Project 55

4-round Waiting for Mr. Perfect Analyze the four-round version of Waiting for Mr. Perfect, whose two round version was discussed in Chapter 26.

CHAPTER 28

Princeton in 1950

The 2001 movie “A Beautiful Mind” focused on game theorist John F. Nash and was a huge success, winning four Academy Awards. It is based on Sylvia Nasar’s unauthorized biography with the same title, published in 1998. There is an amazing passage in the book, which unfortunately was left out of the movie. It is where Nash visited John von Neumann in his office to discuss the main idea and result of the Ph.D. thesis he was writing. At the time Nash was a talented and promising mathematics student at Princeton University and von Neumann was a professor at the Institute for Advanced Study in Princeton, a member of a handful of mathematicians of the highest distinction. Albert Einstein was a colleague. According to the book’s description of the meeting, based on an interview with Harold Kuhn, the visit was short and ended by von Neumann exclaiming: “That’s trivial, you know. That’s just a fixed point theorem” [N1998].

Reading the scene, I see the two men, one eager to impress with his theorem, the other seeming to try to not be impressed. Both men were competitive, more than what may be normal for mathematicians. Although few can answer difficult mathematical questions, many can tell right from wrong answers. For instance, solving an equation may be hard, but checking whether a number is a solution is easy. For this reason, contests are much more frequent in mathematics than in other academic areas. In the 16th century, mathematicians challenged each other, posing questions that they hoped would show their superiority. Nowadays, there are many mathematical contests for students, Putnam Exams, national and international Mathematical Olympiads, and others. For adult mathematicians, there is the game of challenging each other with conjectures, and winning by being the first to prove or disprove one.

Successful mathematicians often seem to be lone wolves. But even wolves need their pack. Mathematicians were physically isolated and able to exchange their ideas only in letters until the end of the 19th century when many mathematical centers emerged, first in Berlin and Paris, and later in Göttingen and European capitals like Budapest, Warsaw, and Vienna. They attracted many world-class mathematicians and students eager to learn from them. There were classes, but many of the exchanges went on during informal activities like walks, private seminars at the home of a professor, or meetings in coffee houses. The main subject of discussion was always mathematics.

Why do mathematicians need their pack? Can’t mathematics be done with paper and pencil (and nowadays, a computer) in isolation? It is nice to learn from the best, but couldn’t one just read their books or papers? Mathematical centers and mathematics conferences are places where mathematicians meet and exchange ideas. In my opinion, they demonstrate that contrary to common perception, mathematics has soft features that require close human interaction. Facts can be learned in books and papers, but the underlying ideas are often not so clearly visible. Mathematicians tend to display only polished results in their publications, not the ideas that led to them. In informal meetings the communication of ideas is the main point. What ideas and approaches have been used? Which have been successful? Which have not? Failed attempts in mathematics are never published. But it is important to know what has been tried unsuccessfully. Math-

ematicians meet to exchange ideas and form and communicate paradigms and research directions that drive research. This is an important way the mathematical community directs a field's development.

Around 1930, through donations and good luck, Princeton changed from a provincial teaching-centered university, where basic knowledge was taught to students, into the mathematical center of the universe [N1998]. Support from the Rockefeller Foundation allowed Princeton University in the mid-20s to create five Research Professorships. These were staffed mostly by Europeans, one of them the very young and very promising John von Neumann. A few years later, another donation helped to create the independent Institute for Advanced Study. Einstein, von Neumann, and two other researcher got positions there. One should also not underestimate the importance of the construction in 1931 of Princeton's Fine Hall, the mathematics building housing a library, comfortable faculty offices furnished with sofas (some with fireplaces!), and, most important, a large common room: excellent facilities for exchanges and for community building [A1988].

It is probably fair to say that game theory started in Princeton in the 1940s. There had been attempts to formalize and analyze games before: for example, an early (1713) analysis and solution of the game *Le Her* discussed in letters between Waldegrave, Montmort, and Nicholas Bernoulli (Chapter 39 of [B2004]), Zermelo's 1913 theorem on chess, which is applicable to all sequential games of perfect information, and papers on poker by Emile Borel. Even von Neumann's main theorem in the field dates back to 1928. But it was in Princeton, where he met the economist Oscar Morgenstern, that von Neumann started the collaboration that resulted in their 1944 book *Theory of Games and Economic Behavior* [VNM1944].

The book got very good reviews and eventually made an enormous impact, but it took a few years before broad research in game theory started. After the publication of the book, von Neumann, as was his habit, turned to other research areas. It was not until 1948 that game theory became a hot topic in Princeton. This was initiated by George Dantzig, who visited von Neumann to discuss linear programming. Von Neumann reacted impatiently to Dantzig's lengthy description of his linear programming setting, but he immediately noticed the connection between it and his theory of two-player zero-sum games [ARD1986], [K2004]. A project investigating the connection between linear programming and game theory was established in Princeton, which included a weekly seminar in game theory. From then on, there was much discussion among faculty and students about game theoretical ideas [K2004], [N1998]. It was this atmosphere that drew Nash towards game theory.

Dantzig's description of his first Princeton visit [ARD1986] seems to indicate that he didn't much enjoy it. But what happened there, a discovery that two areas motivated by different models turn out to be mathematically equivalent, has happened over and over in mathematics and has helped to unify the field. Dantzig must have been grateful that his model turned out to have more applications than he was aware of. We should keep in mind here that von Neumann was not commenting on the work which made Dantzig famous, his invention of the Simplex Method. At that point it was still under construction.

Nash's situation at the time of his meetings with von Neumann was different. His presentation of his main result occasioned von Neumann's "that's trivial" response. But Nash's result extended the range of games considered. Von Neumann's original research was mainly on zero-sum games, and he was initially interested in the case of total competition. Dantzig's linear programs turned out to be models of them: solutions came in solving corresponding linear programs. The von Neumann–Morgenstern book developed a theory for cooperative games, which are games where negotiations are allowed before the game starts and where enforceable contracts can be made. What Nash introduced was a different point of view, a changed perspective: he looked at non-cooperative games, games without enforceable contracts, and maybe even without communication, but in which players don't have completely conflicting interests. This work made game theory applicable to situations outside of parlor or casino games. It changed game theory into a branch of applied mathematics. This is ironic, since von Neumann was interested in applications of mathematics, but Nash was not.

What about the alleged obviousness of Nash's result or his proof? Is it? By all means, no. If you doubt this, try to prove the result yourself. It is described in Chapter 27 on page 204. We know that von Neumann

was a very fast thinker. Maybe he understood Nash's description very quickly and was able to judge it. This might be possible, but it misses the point. Great theorems in mathematics are not necessarily great because they are difficult to prove or understand. Many of the great theorems, like Euclid's theorem that there are infinitely many prime numbers, or Cantor's results about countable and uncountable sets (for example there are as many rational numbers as there are integers, but that there are more real numbers than rational numbers) can be understood by undergraduates. What often makes them great is that they introduce new perspectives, fresh ways to look at something, like Cantor's definition of a cardinal number. Von Neumann did not look carefully at non-zero-sum games, and neither did other mathematicians who were discussing game theory at Princeton at that time. But Nash did, and then he proved his beautiful theorem. So, to summarize, we can say that competition and centers of mathematical thought are important for mathematical progress, as is the ability to ask new questions and create new concepts to answer them. The latter is what Nash did in 1950.

CHAPTER 29

Example: Airport Shuttle

Prerequisites: Chapters 2 and 27.

In Lugano, the hometown of our college, there are two competing shuttle companies that serve the airport in Milano-Malpenso (MXP), which is about 1 hour and 15 minutes away. They have different schedules, but both companies depart from Lugano's railway station. Quite often company A's shuttles are supposed to leave just briefly before company B's. I wondered how many, or actually how few, customers would take the shuttle from company B. Probably only those who arrive at the railway station in the short interval between shuttle A's and shuttle B's departures.

It seems obvious that it matters how the two companies' schedules relate to one another. How should they set up their schedules?

29.1 The Simple Model

In our first model we have two competing airport shuttle companies. To simplify things, we assume that the license granting department requires both to schedule four busses per day, one every four hours, each departure at the start of an hour. Thus the only decision left to Ann and Beth, the owners of the shuttles, is at what hour they should start in the morning, 6:00, 7:00, 8:00, or 9:00. By November 1st they must submit their decisions about next year's schedules. So we may model the situation as a simultaneous game.

29.1.1 To the Airport

Let's concentrate first on the Lugano to airport leg. The shuttles can pick up passengers at the airport for the ride back to the city, which we deal with in Section 29.1.2. We treat both legs together in Section 29.1.3.

Each shuttle's payoff is proportional to the number of passengers it transports. We assume that at the start of each hour between $a:00$ and $b:00$, eight persons have to leave Lugano for the airport, where a and b are parameters we can vary. If no shuttle leaves at that hour, the travelers take a shuttle leaving earlier (if there is one) provided it is not more than three hours earlier—otherwise the traveler seeks another mode of transportation. If two shuttles leave at the same time, they split the passengers equally.

The game is almost but not totally zero-sum. Some travelers may not find a shuttle, either because they have to leave for the airport earlier than the earliest shuttle, or they have to leave more than three hours later than the latest shuttle leaves. They don't take a shuttle.

Let's take $a = 4$ and $b = 23$, so we have passengers needing to leave from 4:00 until 23:00. Assume that Ann's shuttle leaves at 6:00, 10:00, 14:00, and 18:00, and Beth's shuttle leaves at 7:00, 11:00, 15:00, and 19:00. The following table displays the eight passengers' choices at each hour—A for using Ann's shuttle, B for using Beth's, and – for using neither (they have to find another mode of transport).

4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
–	–	A	B	B	B	A	B	B	B	A	B	B	B	A	B	B	B	B	–

So Ann's shuttle leaving at 6:00 will have eight passengers on board, as will Ann's shuttles at 10:00, 14:00, and 18:00. Each of Beth's four shuttles will have 24 passengers except the last, which will transport 32 passengers. Let's assume that shuttles are large, so that they never have to leave passengers behind. Therefore Ann will transport 32 passengers, and Beth 104.

If we do this analysis for all possible pairs of start times for Ann's and Beth's shuttles, we get the following normal form: There is an Excel sheet, [AirportShuttle.xlsx](#), that computes the values automatically.

	6,10,14,18	7,11,15,19	8,12,16,20	9,13,17,21
6,10,14,18	64, 64	32, 104	64, 80	96, 48
7,11,15,19	104, 32	64, 64	32, 104	64, 72
8,12,16,20	80, 64	104, 32	64, 64	32, 96
9,13,17,21	48, 96	72, 64	96, 32	60, 60

Table 29.1. The bimatrix for $a = 4$, $b = 23$, passengers leaving between 4:00 and 23:00

Let us abbreviate the move 6,10,14,18 by M6, and so on. None of the moves is dominated. The maximin move for both players is M9, with a minimum payoff of 48. If both players play their maximin move, they would get a payoff of 60. But this schedule minimizes the total number of passengers transported. The 24 persons who have to leave at 6:00, 7:00, or 8:00 would need other transportation.

The symmetric best response digraph is cyclic: the best response to M6 is M7, the best response to M7 is M8, the best response to M8 is M9, and the best response to M9 is M6. It is best to start one hour later than your opponent. That implies that there is no Nash equilibrium in pure strategies.

Therefore we should look for mixed Nash equilibria. We define three mixed strategies:

- Mix1 is the mixed strategy of using $36/193 \approx 19\%$ of M6, $56/193 \approx 29\%$ of M7, $29/193 \approx 15\%$ of M8, and $72/193 \approx 37\%$ of M9,
- Mix2 is $60/121 \approx 50\%$ of M6, $53/121 \approx 43\%$ of M8, and $8/121 \approx 7\%$ of M9,
- Mix3 is $64/137 \approx 47\%$ of M7, $1/137 \approx 1\%$ of M8, and $72/137 \approx 53\%$ of M9.

Mix2 is a mix of even start times (with a small probability of starting at 9:00) and Mix3 is a mix of odd start times (with a small probability of starting at 8:00).

We claim that Mix1 versus Mix1 is a symmetric Nash equilibrium, and that Mix2 versus Mix3 are also Nash equilibria. Although finding them is at the moment a little beyond what we are supposed to do, we can check for Nashness in a straightforward way. We add the three mixed strategies to the pure strategies matrix, see the matrix on the top of the next page.

The symmetric best response digraph is in Figure 29.1. It turns out that Mix1 versus Mix1 and also Mix2 versus Mix3 are Nash equilibria.

	M6	M7	M8	M9	Mix1	Mix2	Mix3
M6	64, 64	32, 104	64, 80	96, 48	66.7, 72	66.1, 70	65.9, 74.4
M7	104, 32	64, 64	32, 104	64, 72	66.7, 67	69.8, 66.2	63.8, 68.5
M8	80, 64	104, 32	64, 64	32, 96	66.7, 66.7	69.8, 66.1	65.9, 65.9
M9	48, 96	72, 64	96, 32	60, 60	66.7, 63.7	69.8, 65.6	65.9, 61.7
Mix1	72, 66.7	67, 66.7	66.7, 66.7	63.7, 66.7	66.7, 66.7	69.1, 66.7	65.3, 66.7
Mix2	70, 66.1	66.2, 69.8	66.1, 69.8	65.6, 69.8	66.7, 69.1	68, 68	65.9, 69.8
Mix3	74.4, 65.9	68.5, 63.8	65.9, 65.9	61.7, 65.9	66.7, 65.3	69.8, 65.9	64.9, 64.9




Figure 29.1. Best response digraph

Both players get a payoff of $12864/193 \approx 66.65$ in the Mix1 versus Mix1 case. In the other two Nash equilibria there is essentially a mix of M6 and M8 versus a mix of M7 and M9. The player playing Mix2 gets an expected payoff of $9024/137 \approx 65.87$, and the other player gets an expected payoff of $768/11 \approx 69.82$.

The payoffs of the three Nash equilibria are shown in Figure 29.2.

If no communication is allowed between the players, then the game is difficult for both players, since it requires coordination between them to get a large payoff. If communication is allowed before the game starts, then the players could agree to play one of the three Nash equilibria. Which would they agree upon? Without enforceable contracts and side-payments, probably the most likely outcome of the negotiation, the one considered fair by both players, would be the symmetric Nash equilibrium.




Figure 29.2. Payoffs of the three Nash equilibria

Would they obey the pre-game agreement? Yes, of course. Nash equilibria are self-enforcing. If one player deviates, she would make less than or equal to what she would make by sticking to the agreement, provided the other player obeys the agreement. All the mixed Nash equilibria require both players to mix strategies with given probabilities. Since a player couldn't check whether the other kept her promises, wouldn't this invite cheating? No—check the bimatrix to assure yourself that a player deviating and playing a pure strategy cannot improve her expected payoffs.

No pure strategies pair would be the outcome of pre-game negotiations. One player would always be tempted to deviate from such an agreement.

29.1.2 From the Airport

For the ride from the airport to the railway station, we assume that at each hour between $c:00$ and $d:00$, eight persons need a ride. They take the next available shuttle, provided they don't have to wait more than three hours, and shuttles leaving at the same time split the passengers equally.

If we take the same start time $c = a = 4$ and end time $d = b = 23$ for passengers as in the analysis to the airport, we arrive at the following bimatrix. The Excel sheet can be used to compute it.

	7,11,15,19	8,12,16,20	9,13,17,21	10,14,18,22
7,11,15,19	64, 64	104, 32	80, 64	56, 96
8,12,16,20	32, 104	64, 64	104, 32	80, 64
9,13,17,21	64, 80	32, 104	64, 64	104, 32
10,14,18,22	96, 56	64, 80	32, 104	64, 64

This game has only one Nash equilibrium Both players use a mix of $17/33 \approx 52\%$ of 7,11,15,19, $4/11 \approx 36\%$ of 9,13,17,21, and $4/33 \approx 12\%$ of 10,14,18,22, and get a payoff of $2272/33 \approx 68.85$.

29.1.3 Combining Both

When to go to and from the airport are not independent decisions. If a shuttle company has only one shuttle bus, then it cannot depart from the city and the airport at the same hour. If the trip takes more than an hour, then the combination of having a difference of one hour between leaving the city and the airport is impossible. There are only three possibilities:

- 6, 10, 14, 18|8, 12, 16, 20: The bus leaves for the airport at 6, 10, 14, and 18, and goes back at 8, 12, 16, and 20.
- 7, 11, 15, 19|9, 13, 17, 21,
- 8, 12, 16, 20|10, 14, 18, 22.

Let us abbreviate the moves by 6|8, 7|9, and 8|10. Although it is advantageous to depart for the airport exactly one hour after the other shuttle company (since you would collect three rounds of passengers, and the other company only one), it is the opposite for going back. For that, it would be best to leave one hour before the other shuttle. So if you are forced to keep a regular pattern, any advantage you get for departing from the city one hour later than the other company will be balanced by a disadvantage of almost equal weight on the return to the city from the airport.

For instance, if customers need the shuttle from 4:00 to 23:00 both ways, then the payoffs for both shuttle companies are equal in all outcomes:

	6 8	7 9	8 10
6 8	128, 128	136, 136	144, 144
7 9	136, 136	128, 128	136, 136
8 10	144, 144	136, 136	128, 128

We have five Nash equilibria, two pure ones of 6|8 versus 8|10 with payoff of 144 for both, and three mixed ones: 50% 6|8 and 50% 8|10 versus the same strategy, or versus the pure strategy 7|9. In each of the mixed Nash equilibria, the payoffs for both players are 136.

29.2 Impatient Customers

In this variant we make the assumption that the number of passengers who use a shuttle decreases if they would have to wait too long at the airport. We suppose that out of the eight persons who have to leave Lugano at a given hour, one chooses other transportation (like a taxi, or letting a friend drive him or her) if he or she has to leave one hour earlier, and three (respectively six) choose other transportation if they would have to leave two (respectively three) hours earlier. Remember that nobody having to leave four or more hours earlier would use the shuttle. The same assumptions are made for the trip from MXP to Lugano. One, three, or six of the eight passengers would seek other transportation if they would have to wait one hour, two hours, or three hours.

Assume that $a = 6$ and $b = 20$ so that the earliest passengers want to go to the airport at 6:00 and the latest at 20:00, and $c = 7$ and $d = 21$ so that the earliest passengers want to leave the airport at 7, and the latest at 21:00. We can use the Excel sheet to compute the bimatrix:

	6 8	7 9	8 10
6 8	83.5, 83.5	107, 107	120, 105
7 9	107, 107	83.5, 83.5	112, 92
8 10	105, 120	92, 112	77, 77

Although for customers the combination 6|8 versus 8|10 would be best, because 225 of 240 possible passengers would be transported on shuttles, the combination is not a Nash equilibrium. The player playing 8|10 would want to change to 7|9. There are five Nash equilibria, two pure ones 6|8” versus 7|9 with a payoff of 107 for both, and three mixed. The two pure ones have also the highest number of passengers transported (214) of the Nash equilibria. Maybe the licensing authority should suggest this to both companies. Although the payoffs are equal, the equilibria are not symmetric, so pre-game negotiations are necessary about who will start at 6:00 and who at 7:00.

Exercises

- 1. Try to get data from your local airport about how many passengers leave and arrive each hour. You could look through the departure and arrival schedules, or ask the local shuttle service whether it has data, or count yourself. Use the Excel sheet to analyze the game.
- 2. If your airport has two shuttles, look up their schedules and find out whether they form a Nash equilibrium, either with an assumption on the passenger distribution or with the data in the previous question.

CHAPTER 30

Example: Election II

Prerequisites: Chapters 2, 7, and 27.

This chapter is a continuation of Chapter 7, where we described versions of the games $\text{ELECTION}(c, d, e|c_1, d_1, e_1|a, b)$. We did not discuss the many versions that don't have pure Nash equilibria, but now mixed strategies allow us to take a different tack. Since we can model these games as two-person zero-sum games, von Neumann's and Robinson's theorems apply, and the mixed Nash equilibria are meaningful. We will use the Excel sheet [Election2.xlsm](#), where the payoff matrix is calculated automatically, and where we implement Brown's fictitious play method to approximate mixed Nash equilibria.

We lay a foundation for this chapter with a few assumptions about campaigns, most obvious and supported by empirical evidence collected by the organization FairVote [[Fn.d.](#)], that monitored the close 2004 Bush and Kerry campaigns for President. We will then analyze some simple examples to see whether the assumptions are supported by game theory.

30.1 Left Over from Election I

Before we state our assumptions, we will complete the analysis of the second example $\text{ELECTION}(7, 8, 13|-1, -1, 1|4, 4)$ from Chapter 7, which we couldn't finish since the game doesn't have a Nash equilibrium in pure strategies. After repeatedly eliminating weakly dominated moves, we arrived at the matrix:

	0,0,4	1,0,3	0,1,3	1,2,1
1,0,3	-1	-1	1	-1
0,2,2	-1	1	-1	-1
0,3,1	-1	-1	-1	1
2,2,0	1	-1	-1	-1

Let the mixed strategy Mix_A for Ann denote the mix of the moves (1, 0, 3), (0, 2, 2), (0, 3, 1), and (2, 2, 0), each with equal probability 25%. Take as the mixed strategy Mix_B for Beth the equal-probability mix of her moves (0, 0, 4), (1, 0, 3), (0, 1, 3), and (1, 2, 1). Then Ann's Mix_A achieves an expected payoff of $-\frac{1}{2}$ against all of Beth's remaining moves. Therefore Mix_A will achieve a payoff of $-\frac{1}{2}$ against every mix of Beth's four moves. In the same way, Beth's Mix_B achieves an expected payoff of $\frac{1}{2}$ against all of Ann's remaining moves, and therefore also against every mix of them. Thus Mix_A and Mix_B are maximin mixes for Ann and Beth and form a Nash equilibrium by von Neumann's Theorem.

Since Beth's moves (1, 2, 1) and (2, 2, 0) perform identically against Ann's remaining moves, Mix_A forms a Nash equilibrium with any of Beth's mixes of 25% of (0, 0, 4), 25% of (1, 0, 3), 25% of (0, 1, 3), $x\%$ of (1, 2, 1), and $(25 - x)\%$ of (2, 2, 0).

30.2 More Effort into Large Districts

The report [Fn.d.] presents data for money spent on ads and number of visits per state by the candidates for the last thirty-seven days of presidential campaigns. The data can also be found in the Excel sheet [ElectionCampaign2004.xlsx](#). As in the 2012 presidential election, many states didn't receive any visits or money for ads, because they were considered to be already lost or won. Only swing states, also called battleground states attracted the campaigns' attention. It would be interesting to find out which states were considered swing states at that time—Florida, where the result was not determined for days in the 2000 election was certainly among them. But since we don't know the pre-election perceptions of the parties, we take a post-election point of view and just look at the states where the popular vote for Bush and Kerry differed by at most 5%. Figure 30.1 plots the number of visits by the candidates to the states against its number of electoral votes, filled black squares for the Democrats and empty diamonds for the Republicans. Figure 30.1 also shows regression lines, the straight lines closest to the data points. Since they (almost) go through the origin, the relationship between the variables is close to proportional—twice the number of votes attracts twice the number of visits. The number of electoral votes is a good predictor of the number of visits for the swing states.




Figure 30.1. Visits versus number of electoral votes for Kerry and Edwards (black squares) and Bush and Cheney (empty diamonds) in different states in the 2004 US Presidential Election

Let us try to confirm this pattern in our small example, where the three districts have different sizes and the sum of votes in the two smaller districts exceeds the number of votes in the largest district. This was the case we called 7-8-13 in Chapter 7. Although a player winning any two districts wins the presidency, if one district has a tie, the candidate winning the larger district wins the presidency.

Using the Excel sheet [Election2.xlsx](#) and running Brown's fictitious play 5000 times we find mixed strategies for the symmetric games $\text{ELECTION}(7, 8, 13|0, 0, 0|3, 3)$, $\text{ELECTION}(7, 8, 13|0, 0, 0|4, 4)$,

and $\text{ELECTION}(7, 8, 13|0, 0, 0|5, 5)$. In $\text{ELECTION}(7, 8, 13|0, 0, 0|3, 3)$, both Ann and Beth use 20% of each of $(0, 0, 3)$, $(1, 0, 2)$, $(0, 1, 2)$, $(1, 1, 1)$, and $(1, 2, 0)$. Therefore, with 20% probability a player puts into the smallest district C 0, 1, 0, 1, and 1 resources. The expected value of resources put into district C by each player is $20\% \cdot 0 + 20\% \cdot 1 + 20\% \cdot 0 + 20\% \cdot 1 + 20\% \cdot 1 = 0.6$. Similarly, each player puts $20\% \cdot 0 + 20\% \cdot 0 + 20\% \cdot 1 + 20\% \cdot 1 + 20\% \cdot 2 = 0.8$ resource into district D, and $20\% \cdot 3 + 20\% \cdot 2 + 20\% \cdot 2 + 20\% \cdot 1 + 20\% \cdot 0 = 1.6$ resources into the largest district E. Therefore both put on average 20% of their resources into district C, 27% into district D, and 53% into district E.

This analysis may oversimplify matters. It is not enough to tell the players how frequently resources should be put into the districts. One can find mixed strategies with the same percentages that give different payoffs.

Doing the same for the other two games, we find similar values in $\text{ELECTION}(7, 8, 13|0, 0, 0|4, 4)$: Both players on average put 23% of their resources into district C, 27% into district D, and 51% into district E. For $\text{ELECTION}(7, 8, 13|0, 0, 0|5, 5)$ the numbers are closer together: 25%, 30%, and 45%. Having more resources and more time weakens the emphasis on the large district.

Do we also observe proportionality between effort and number of votes? Yes, in all three games, but the question may be ill-posed, since the same payoff matrices, the same solutions, and the same distributions of effort into the districts occurs when the districts have different number of votes, and together the two smaller districts have more votes than the largest one. So, we get the same effort distributions for 2, 12, 13 votes, 11, 12, 13 votes, or 7, 8, 13 votes, and the values cannot be proportional in all these cases.

30.3 Defend Where Ahead or Attack Where Weak?

Now let's look at asymmetric games where Ann has an advantage in one district and Beth in another. Should a player put more resources into the district where she has an advantage, or should she try to turn the district where she is behind? We use the same method as in the previous section, finding first mixed Nash equilibria and then adding the expected resources in the districts for the mixed strategies.

In $\text{ELECTION}(7, 8, 13|0, -1, 1|3, 3)$ Ann is one resource behind in district D and leads by one resource in district E. One mixed Nash equilibrium found by Brown's fictitious play in the Excel sheet consists of about

- 9% of $(0, 0, 3)$, 18% of $(1, 0, 2)$, 18% of $(0, 1, 2)$, 18% of $(0, 2, 1)$, and 37% of $(1, 2, 0)$ for Ann, and
- 26% of $(0, 0, 3)$, 13% of $(0, 1, 2)$, 28% of $(0, 1, 2)$, 8% of $(1, 1, 1)$, 19% of $(2, 1, 0)$, and 5% of $(1, 2, 0)$ for Beth.

On average, Ann puts 18%, 42%, and 39% of her resources into districts C, D, and E, and Beth puts 22%, 22%, and 56% of her resources into them. The expected payoff for Ann is about 0.091. Thus Ann has an advantage, which is obvious since the district where she leads is larger than the district where Beth is ahead. Ann puts the most effort into district D, where she is behind, and Beth puts the most effort into district E, where she is behind. The effect is more visible for Beth, who has a disadvantage.

$\text{ELECTION}(7, 8, 13|0, -1, 1|4, 4)$ confirms this picture: Ann puts 26%, 40%, and 34% of her resources into districts C, D, and E; Beth puts in 23%, 17%, and 60%. The expected payoff for Ann is 0.021. For a higher number of available resources, the patterns seems to shift, as in $\text{ELECTION}(7, 8, 13|0, -1, 1|5, 5)$ the numbers are 24%, 38%, 38% for Ann, and 23%, 24%, 53% for Beth. Ann seems to shift from attacking in district D to defending her advantage in district E. Also, Ann's advantage diminishes further to a value of only 0.008.

Based on these cases, we may suggest that Ann and Beth should on average put more effort into the districts where they have a disadvantage, at least if the number of resources still available is small. This recommendation seems to be more valid for the player who has lower expectations.

What about the real world? Look at the largest swing states Florida, Pennsylvania, Ohio, Michigan, and at Wisconsin and Iowa, which got more visits than the remaining states. Bush won in Florida, Ohio, and

Iowa, Kerry won in Pennsylvania, Michigan, and Wisconsin. In the states the candidate who invested less effort won. If we assume that the states had been considered swing states before, an assumption certainly true for Florida and supported by the fact that the candidates put so much effort into them, we can draw two conclusions. Either people dislike politicians so much that more effort means fewer votes, an assumption which sometimes has some merit but overall still seems too cynical, or the candidate investing more was initially behind, but didn't quite make it. We have a confirmation that candidates put more effort into states where they are a little behind.

As an exercise, use Excel sheet [ElectionCampaign2004.xlsx](#) to check whether the money spent on ads confirms this as well. (See Exercise 1 below.)

30.4 Is Larger Better?

ELECTION(7, 8, 13| − 1, 1, 0|3, 3) has many mixed Nash equilibria. The payoff matrix is

	0,0,3	1,0,2	0,1,2	2,0,1	1,1,1	0,2,1	3,0,0	2,1,0	6 1,2,0	0,3,0
0,0,3	1	1	1	1	1	−1	1	1	−1	−1
1,0,2	−1	1	0	1	1	1	1	1	−1	1
0,1,2	−1	1	1	1	1	1	1	1	1	−1
2,0,1	1	−1	−1	1	0	−1	1	1	1	1
1,1,1	−1	−1	−1	1	1	0	1	1	1	1
0,2,1	−1	−1	−1	1	1	1	1	1	1	1
3,0,0	1	1	−1	−1	−1	−1	1	0	−1	−1
2,1,0	1	−1	1	−1	−1	−1	1	1	0	−1
1,2,0	−1	−1	−1	−1	−1	−1	1	1	1	0
0,3,0	−1	−1	−1	−1	−1	−1	1	1	1	1

Take a mix of x of (0, 1, 2), x of (2, 0, 1), y of (0, 0, 3), and y of (1, 0, 2) for Ann, where $2x + 2y = 1$. It achieves these payoffs against the pure moves of Beth:

0,0,3	1,0,2	0,1,2	2,0,1	1,1,1	0,2,1	3,0,0	2,1,0	6 1,2,0	0,3,0
0	2y	y	2x + 2y	x + 2y	0	2x + 2y	2x + 2y	2x − 2y	0

These values are nonnegative for $x \geq y$. Therefore Ann can achieve a payoff of at least 0 using this mix against Beth. For Beth, consider a 50%-50% mix of (0, 0, 3) and (0, 2, 1). It guarantees a payoff of 0 for Beth against any of Ann's pure strategies. Therefore every $x - y$ mix with $x \geq y$ for Ann versus the 50%-50% mix for Beth forms a mixed Nash equilibrium. The expected payoff is 0, even though the district where Ann has an advantage is larger than the district where Beth leads. Hence the answer to the question for this section is no; leading in the largest district and being behind in a smaller one is not always advantageous.

Ann's $x - y$ mix puts $(2x + y)/3$ into district C, $x/3$ into district D, and $(3x + 5y)/3$ into district E. The numbers vary from 25%, 8%, 67% for $x = 25\%$ to 33%, 17%, 50% for $x = 50\%$. Beth's mix puts nothing into district C, 33% into district D, and 67% into district E. Thus both should put the most effort into the largest district, and more effort into the district where they are behind than into the district where they are leading, as in most of the previous examples.

30.5 ELECTION(7, 8, 13| − 1, −1, 2|x, x)

Is large and large better than middle and small? Is Ann having a larger, 2-resource lead in the large district better off than Beth who leads in both smaller districts by one resource each? One might think so, but our models show that this does not have to be the case. The first two games give an expected value of 0.

- In $\text{ELECTION}(7, 8, 13 | -1, -1, 2 | 3, 3)$ a mixed Nash equilibrium is Ann's equal mix of $(0, 2, 1)$ and $(2, 1, 0)$ versus Beth's equal mix of $(0, 0, 3)$ and $(1, 2, 0)$, with an expected payoff of 0 for both.
- In $\text{ELECTION}(7, 8, 13 | -1, -1, 2 | 4, 4)$ a mixed Nash equilibrium is Ann's equal mix of $(0, 2, 2)$, $(3, 0, 1)$, $(0, 3, 1)$, $(2, 2, 0)$ versus Beth's mix of 30% of $(0, 0, 4)$, 20% of $(1, 0, 3)$, 30% of $(2, 2, 0)$, and 20% of $(1, 3, 0)$.

The case with five resources differs. See the exercises.

Exercises

1. Use the data in the Excel sheet [ElectionCampaign2004.xlsx](#) to graph the money spent on ads for both candidates against the number of electoral votes for swing states. Are these numbers roughly proportional? Can the pattern that candidates win who put less effort into the state be confirmed for this data?
2. Who has an advantage in $\text{ELECTION}(7, 8, 13 | -1, -1, 2 | 5, 5)$?
3. Analyze $\text{ELECTION}(7, 8, 13 | -1, 2, -1 | 3, 3)$.
4. Analyze $\text{ELECTION}(7, 8, 13 | 0, 0, 0 | 4, 4)$. Compare the allocation of resources to C, D, and E in the mixed Nash equilibrium with the values given in Section 30.2.
5. Analyze $\text{ELECTION}(7, 8, 13 | 0, 0, 0 | 5, 5)$. Compare the allocation of resources to C, D, and E in the mixed Nash equilibrium with the values given in Section 30.2.
6. Analyze $\text{ELECTION}(7, 8, 13 | 0, -1, 1 | 4, 4)$. Compare the allocation of resources to C, D, and E in the mixed Nash equilibrium with the values given in Section 30.3.
7. Analyze $\text{ELECTION}(7, 8, 13 | 0, -1, 1 | 5, 5)$. Compare the allocation of resources to C, D, and E in the mixed Nash equilibrium with the values given in Section 30.3.
8. What happens in $\text{ELECTION}(7, 8, 13 | -1, 3 - 1 | 3, 3)$? How many resources do the players send into the districts on average? How do things change in $\text{ELECTION}(7, 8, 13 | -1, 3 - 1 | 4, 4)$? What happens in $\text{ELECTION}(7, 8, 13 | -1, 3 - 1 | 5, 5)$?

CHAPTER 31

Example: VNM POKER(2, r , m , n)

Prerequisites: Chapter 25 and all previous theory chapters.

In this chapter we use mixed strategies to look at versions of VNM POKER(2, r , m , n). It gives exercise in working with parameters, manipulating algebraic expressions, and graphing equations.

We play VNM POKER(2, r , m , n) with r cards of value 1 and r cards of value 2, with initial bets of m and raised bets of n . We assume $r \geq 2$. The extensive form of VNM POKER(2, 4, 2, 3) is shown in Figure 31.1.

Class Activity In the applet [VNMPoker2](#), select values of m and n and a computer opponent (Chingis, Grandpa, or Lucky) and play 30 rounds. Try to win money overall—it is possible. If you didn't succeed, try for another 30 rounds.

In our case $S = 2$, the pure strategies for Ann are the prudent CC (always check), the reasonable CR (check with a 1 and raise with a 2), the silly looking RC (raise with a 1 and check with a 2) and the aggressive




Figure 31.1. VNM POKER(2, 4, 2, 3)

	FF	FC	CF	CC
CC	0	0	0	0
CR	$\frac{m(r-1)}{4r-2}$	0	$\frac{nr-m}{4r-2}$	$\frac{(n-m)r}{4r-2}$
RC	$\frac{(3r-1)m}{4r-2}$	$\frac{(2r-1)m-rn}{4r-2}$	$\frac{2mr}{4r-2}$	$\frac{(m-n)r}{4r-2}$
RR	m	$\frac{(2r-1)m-rn}{4r-2}$	$\frac{(2r-1)m+rn}{4r-2}$	0

Table 31.1. Normal form of VNM POKER(2, r, m, n)

RR (raise in any case). Beth's pure strategies are the prudent FF (fold in any case), the reasonable FC (fold with a 1, call with a 2), the counterintuitive CF (call with a 1, fold with a 2), and the aggressive CC (always call). The normal form of Ann's expectations is seen in Table 31.1

The expressions are obtained as follows. For a pair of strategies, let $A_{x,y}$ be Ann's payoff provided both players stick to their chosen strategies, Ann gets a card of value x and Beth a card of value y . For instance, if Ann plays the aggressive RR and Beth plays the reasonable FC, then $A_{1,1} = m$, $A_{1,2} = -n$, $A_{2,1} = m$, and $A_{2,2} = 0$. Then Ann's expected payoff for the pair of strategies is

$$p_{XX} \cdot A_{1,1} + p_{XY} \cdot A_{1,2} + p_{XY} \cdot A_{2,1} + p_{XX} \cdot A_{2,2},$$

with the probabilities for equal cards $p_{XX} = \frac{r-1}{4r-2}$ or for different cards $p_{XY} = \frac{r}{4r-2}$, as explained earlier. For the choice of strategies, RR versus FC, we get

$$p_{XX} \cdot m + p_{XY} \cdot (-n) + p_{XY} \cdot m + p_{XX} \cdot 0$$

as Ann's payoff, which is

$$\frac{r-1}{4r-2} \cdot m + \frac{r}{4r-2} \cdot (-n) + \frac{r}{4r-2} \cdot m + \frac{r-1}{4r-2} \cdot 0 = \frac{(r-1)m - rn + rm}{4r-2} = \frac{(2r-1)m - rn}{4r-2}.$$

As we noted in Chapter 25, because Ann's strategy CC is weakly dominated by CR, and RC is weakly dominated by RR, both can be deleted. Similarly, Beth's strategy FF is weakly dominated by FC, and CF is weakly dominated by CC. The resulting normal form is

	FC	CC
CR	0	$\frac{(n-m)r}{4r-2}$
RR	$\frac{(2r-1)m-rn}{4r-2}$	0

Table 31.2.

31.1 The Case $\frac{n}{m} \geq 2 - \frac{1}{r}$

The entry $\frac{(n-m)r}{4r-2}$ is always positive. If $\frac{(2r-1)m-rn}{4r-2}$ is not positive, then Ann's strategy CR weakly dominates RR, and Beth's strategy FC weakly dominates CC. Therefore there is an equilibrium in pure strategies, CR versus FC, with an expected payoff of 0 for Ann.

The value $\frac{(2r-1)m-rn}{4r-2}$ is non-positive if $rn \geq (2r-1)m$, i.e., if $\frac{n}{m} \geq \frac{2r-1}{r} = 2 - \frac{1}{r}$, i.e., if n is considerably larger than m . The values $2 - \frac{1}{r}$ equal $\frac{3}{2}$, $\frac{5}{3}$, and $\frac{7}{8}$ for $r = 2, 3, 4$. Even for large r , the value is always smaller than 2. If raising means doubling, the analysis is finished for every r : aggressive play, or bluffing, will not occur, since it is too expensive.

31.2 Best Responses

How would you play against a player who doesn't necessarily play optimally, but mixes the two non-dominated pure strategies arbitrarily?

Assume Beth mixes her pure strategies: she plays FC with probability q , and CC with probability $1 - q$. That means that Beth always calls when holding a card of value 2, and folds with probability q when holding a card of value 1. When would Ann respond to the mix with CR and when with RR?

Ann's payoff when playing CR is $\frac{(1-q)(n-m)r}{4r-2}$. This is larger or equal to Ann's payoff when playing RR, which is $\frac{q((2r-1)m-rn)}{4r-2}$, if

$$\frac{(1-q)(n-m)r}{4r-2} \geq \frac{q((2r-1)m-rn)}{4r-2}.$$

We solve this inequality for q . Because $4r - 2 > 0$,

$$(1-q)(n-m)r \geq q((2r-1)m-rn)$$

or

$$(n-m)r - q(n-m)r \geq q((2r-1)m-rn)$$

so

$$(n-m)r \geq q[(2r-1)m-rn + (n-m)r]$$

which simplifies to

$$(n-m)r \geq q[(r-1)m].$$

We divide by the positive value $(r-1)m$ and get

$$\frac{(n-m)r}{(r-1)m} \geq q.$$

Denote $\frac{(n-m)r}{(r-1)m}$ by q^* . Ann should play the reasonable, bluff-free strategy CR when Beth plays FC with probability less than q^* , i.e., if Beth plays too aggressively. Ann should otherwise play the aggressive strategy RR.

In the same way we compute Beth's best response to Ann's mixed strategy of playing CR with probability p , and playing RR with probability $1 - p$ (which translates into checking with probability p when holding a card of value 1, and raising with a card of value 2). When Beth plays FC, Beth's expected payoff is $-\frac{(1-p)((2r-1)m-rn)}{4r-2}$ and when she plays CC, it is $-\frac{p(n-m)r}{4r-2}$. Both values are the same for $p = \frac{(2r-1)m-rn}{(r-1)m}$. Call this value p^* . Beth's payoff when she plays FC is smaller than her payoff when she plays CC provided $p < p^*$. Beth should play CC when Ann plays CR with probability less than p^* , when Ann plays too aggressively and bluffs too much. Conversely, Beth should lamely play FC when Ann plays CR more often, that is, if Ann plays too lamely.

If she is clever, Ann will play opposite to the observed behavior of Beth. Beth, on the other hand, should mimic the play she observes in Ann, whether it is aggressive or cautious, .

31.3 Reading the Opponent (optional)

Class Activity Repeat the class activity at the beginning of the chapter, keeping in mind our analysis of best responses from the previous section.

The key to effective play is figuring out how the computer player plays by observing its play. Assume that Ann always raises and Beth always calls with a card of value 2—without those assumptions the analysis would be more difficult.

If players randomize their strategies, we cannot determine mixed strategy play probabilities after only a few rounds. After many rounds, though, the Law of Large Numbers tells us that the observed relative frequencies should be close to the probabilities in the mixed strategy. A problem is that we cannot observe all decisions—we don't see all of the other player's decisions. If a player folds, then no player sees the other's card, not even after the game is over.

What Beth can observe about Ann's play is how often Ann checks and raises, say 30% and 70%, for example. We assume that Ann checks only with a card of value 1. Since in the long run, in about 50% of all rounds Ann has a 1, but she checks in only about 30% of the rounds, we can conclude that Ann raises with a 1, in about 20% of the rounds. Therefore in $\frac{20\%}{50\%} = 40\%$ of the cases where Ann has a 1 she raises. In other words, Ann plays a mix of 60% CR and 40% RR.

Let us now discuss how Ann can observe Beth's strategy. We look only at the cases where Ann raises with a 2. Let's say Beth calls in 70% of those cases, and when she calls she displays 39% of 1s and 61% of 2s. Among the cases considered, Beth calls with a 1 with frequency $70\% \cdot 39\% \approx 27\%$, and with a 2 with frequency $70\% \cdot 61\% \approx 43\%$. 43% is about $3/7$, the fraction of cases where Beth holds a 2, given that Ann holds a 2, which confirms our assumption that Beth never folds with a 2. The only other option is to fold with a 1, which occurs in 30% of the cases. Therefore, when Beth has a 1 she calls in $\frac{27\%}{27\%+30\%} \approx 47\%$ of the cases. Thus in our example Beth plays a mix of 53% FC and 47% CC.

31.4 Mixed Nash Equilibrium for $\frac{n}{m} \leq 2 - \frac{1}{r}$

If $\frac{n}{m} \leq 2 - \frac{1}{r}$, since there is no pure Nash equilibrium, there must be a mixed Nash equilibrium when Ann plays CR with probability p and RR with probability $1 - p$ and Beth plays FC with probability q and CC with probability $1 - q$ (as usual $0 < p < 1$ and $0 < q < 1$). Since, as a Nash equilibrium, each mixed strategy is a best response to the other, and therefore each of the pure strategies CR and CC, respectively FC and CC, is a best response to them, we conclude, as we did above, that $p = p^* = \frac{(2r-1)m-rn}{(r-1)m}$ and $q = q^* = \frac{(n-m)r}{(r-1)m}$.

They can be expressed in terms of the ratio $\frac{n}{m}$ as:

$$p^* = \frac{2r-1-r\frac{n}{m}}{r-1} \quad \text{and} \quad q^* = \frac{(\frac{n}{m}-1)r}{r-1}.$$

This means that how the players mix depends only on the ratio $\frac{n}{m}$. If we double both ante m and raised bet n , the players would not change their play.

We now compute the value of the game, Ann's expected payoff when Ann and Beth use their Nash equilibrium mixed strategies.

$$V = p^* \cdot \frac{(n-m)r}{4r-2} + (1-p^*) \cdot 0 = \frac{2r-1-r\frac{n}{m}}{r-1} \cdot \frac{(n-m)r}{4r-2} = m \cdot \frac{(2r-1-r\frac{n}{m})(\frac{n}{m}-1)r}{(r-1)(4r-2)}$$

The game's value could also be computed as

$$p^* \cdot 0 + (1-p^*) \cdot \frac{((2r-1)m-rn)}{4r-2} \quad \text{or as} \quad q^* \cdot 0 + (1-q^*) \cdot \frac{(n-m)r}{4r-2}$$

or as

$$q^* \cdot \frac{((2r-1)m-rn)}{4r-2} + (1-q^*) \cdot 0.$$

Because V is the product of the initial bet m and an expression in $\frac{n}{m}$ and r , if we keep $\frac{n}{m}$ and r fixed (and change n proportionally with m) v is proportional to m . So the value per ante $\frac{V}{m}$ depends only on r and the ratio $\frac{n}{m}$:

$$\frac{V}{m} = \frac{(2r-1-r\frac{n}{m})(\frac{n}{m}-1)r}{(r-1)(4r-2)}.$$




Figure 31.2. $\frac{V}{m}$ in terms of $\frac{n}{m}$ and r

This relation is visualized in the graph in Figure 31.2. The x -axis displays r (from 2 to 8), the y -axis displays $\frac{n}{m}$ (from 1 to 2), and the color expresses the value $\frac{V}{m}$, according to the formula. The curves, called *contour lines*, connect points with identical $\frac{V}{m}$. For three of them, the value $\frac{V}{m}$ is also given. The area on the top in purple is the area with $V = 0$ where the pure strategies CR and FC apply, and the curve below it is the curve $\frac{n}{m} = 2 - \frac{1}{r}$. Because r is a positive integer, the value 1.5 for r does not make sense, but the graph expresses a lot of information about the game.

31.5 Small Changes in the Parameters

When playing the game, suppose somebody proposes to increase n slightly. To whose advantage would it be, Ann's or Beth's? Or suppose we increase the number r of duplicates. Which player will profit?

The questions can be answered by looking at the graph in Figure 31.2. Increasing r means we move to the right on the graph. The ratio $\frac{V}{m}$ increases slightly, where V is the value of the game. Thus playing with larger decks of cards is advantageous for Ann.

Changing n is more complicated. Increasing n increases $\frac{n}{m}$ if m is fixed. That means we move vertically on the graph in Figure 31.2. Starting at the r -axis, $\frac{V}{m}$ increases until we meet the ridge, which is visualized by the dashed line in the graph. (The ridge curve has the equation $\frac{n}{m} = \frac{3}{2} - \frac{1}{2r}$. You can show this by using calculus.) Therefore the answer is that a slight increase in n causes Ann's expected payoff to increase provided we are south of the ridge. North of it Ann's expected payoff decreases if n increases slightly.

To answer the question of how the Nash equilibrium mixes would change if we increase r or n , let us graph the values p^* and q^* in terms of r and $\frac{n}{m}$. Remember that p^* and q^* are respectively the probabilities of Ann checking and Beth folding when holding a card of value 1. In the graphs in Figure 31.3 the colors correspond to the values of p^* and q^* , purple being 0 (aggressive play), greenish being around 0.5, and red indicating 1. The contour curves are combinations with constant p^* and q^* .

We can see that south of the curve $\frac{n}{m} = 2 - \frac{1}{r}$ the mixes of Ann and Beth complement each other. That is, when Ann raises a lot with cards of value 1, Beth calls little, and vice versa. Just north of the curve Ann always checks with a 1 and Beth folds with a 1. There the stakes are too high for bluffing.

From the graphs it follows that increasing r means that p^* increases slightly while q^* decreases—Ann bluffs less but Beth calls more with a 1. Slightly increasing n results in the opposite—a decrease of p^* and an increase of q^* . In any case, small changes of the parameters result in small changes of p^* and q^* , except if we cross the curve $\frac{n}{m} = 2 - \frac{1}{r}$. In that case p^* changes from 1 to 0 or conversely. A slight cause with dramatic implications.




Figure 31.3. p^* and q^* in terms of $\frac{n}{m}$ and r

Exercises

1. Confirm the values given for the strategies CR versus FC and for RC versus FC in Table 31.2.
2. Confirm the values given for the strategies CR versus CC and for RR versus CF in Table 31.2.
3. Find the mixed strategy of Grandpa in the applet [VNMPoker2](#) by playing 50 rounds and using the method of Section 31.3.
4. Find the strategies of Chingis and Lucky in the applet [VNMPoker2](#) by playing 50 rounds and using the method of Section 31.3.

CHAPTER 32

Theory 9: Behavioral Strategies

A pure strategy is like instructions written in a notebook that an overprotective father gives his child before the child goes to middle school in a big city. They should describe clearly and precisely what to do in each every situation the child might encounter. Many will not occur, because of chance or because the other persons don't make the decisions that lead to them, or because the child makes decisions that prevent them from occurring. The child can avoid some situations but not others. The strategy prepares the child for every possible situation.

Continuing the metaphor, we can think of a mixed strategy as a set of notebooks to shuffle before leaving for school, choosing one to follow according to a probability distribution. In one play (one day at school), the child would take only one notebook to school, but which one would be unpredictable. A single random choice suffices for behavior in all situations to look random.

Behavioral strategies differ from mixed strategies. Instructions are written in a notebook and cover every situation. Instead of giving definite instructions, the notebook lists all options and assigns probabilities to them. As with mixed strategies, the other players do not know what strategies the child will play. The child would use a random device, like a pair of dice, to determine what option to choose. As with mixed strategies, the behavior of the child is unpredictable, but unlike mixed strategies the child consults the random device repeatedly.

If you think that behavioral strategies approximate the way humans play a game, be careful! Don't confuse behavioral strategies with a player waiting to see what happens and then deciding what move to play. There is nothing local about behavioral strategies. The child's decisions depend only on the probabilities given in the notebook, and they, although each refers to a single situation, are determined with the whole game in mind, globally and in advance. There is nothing spontaneous in a behavioral strategy. Once a player has decided on a pure strategy, a mixed strategy, or a behavioral strategy there is no more free will.

While mixed strategies refer to the normal form of the game (as in Chapter 31), behavioral strategies refer to the game's description in extensive form. A behavioral strategy for a player describes for each of her information sets a probability distribution—probabilities for the options at that information set that sum to 1.

Example 1 Take as an example $VNMPOKER(2, 4, 3, 2)$, which we analyzed in Chapter 31. The extensive form of the game is in Figure 31.1. Ann has two information sets, which are described in the figure by the dashed curves, as does Beth. Ann's information sets depend on the value of Ann's card (1 or 2) and Beth's depend on Beth's card. An example for a behavioral strategy for Ann would be for her to raise 30% of the time and check otherwise when she holds a 1, and when she holds a 2 to raise in 80% of the time and check otherwise.

32.1 Behavioral versus Mixed Strategies

It seems natural to use mixed strategies when the game is given in normal form, and behavioral strategies when it is given in extensive form. However, as we have seen in Chapter 24, extensive form games can be translated into normal form, so which one should we use? Often, behavioral strategies are easier to describe, since there are usually more mixed strategies than behavioral strategies. Easy description is an advantage, but could it be a disadvantage? If there are more mixed strategies, isn't it possible that there are very good mixed strategies that have no counterpart in behavioral strategies? We will see that this is not the case.

Before we do that, let me justify the claim on the numbers. Assume a player has n information sets, each with two options. Then the player has 2^n pure strategies. A mixed strategy is described by assigning a probability for each of them, so it is described by 2^n numbers. A behavioral strategy, on the other hand, is described by only $2n$ probabilities: two probabilities for the two moves in each of the n information sets, a smaller number than 2^n . Although there may be many fewer reduced pure strategies than pure strategies, there are usually fewer behavioral strategies than mixed strategies.

Example 2 To illustrate the subtleties, we use a more complex example than VNMPOKER(2, 4, 3, 2). We will use the game whose extensive form as game tree is in Figure 32.1.




Figure 32.1. Another example

Ann has three information sets. The first is the starting position, where she chooses between A_1 and A_2 . The second consists of two vertices, where she has the options A_3 and A_4 , and the third has two vertices with options A_5 and A_6 . Beth has three information sets. The first consists of two vertices and has options B_1 and B_2 , the second is a single vertex set with options B_3 and B_4 , and the third has two vertices and options B_5 and B_6 . There are also two random moves.

The reduced strategies for Ann are $A_1 A_3 \bullet$, $A_1 A_4 A_5$, $A_1 A_4 A_6$, and $A_2 \bullet \bullet$. The symbol “ $\bullet$ ” is a placeholder meaning that any option available could be used (since in practice the information set will not occur when playing this strategy). Beth's reduced strategies are $B_1 \bullet B_5$, $B_1 \bullet B_6$, $B_2 B_3 B_5$, $B_2 B_3 B_6$, $B_2 B_4 B_5$, and $B_2 B_4 B_6$.

32.1.1 Calculating Mixed Strategies from Behavioral Strategies

Every behavioral strategy can be modeled by a mixed strategy: we need only define the probabilities of each player's pure reduced strategies. Recall that a pure reduced strategy is a sequence of move choices for a player's information sets. Each of the selected move choices has a probability in the behavioral strategy. Then the probability that the player will decide the same as in the pure reduced strategy in all situations is the product of the probabilities. For an unreachable information set, as indicated by the symbol " $\bullet$ ", we multiply by 1. We take the product of the probabilities because we assume that decisions about moves are independent choices.

Example 2 continued Assume that in the first information set Ann chooses A_1 with probability $7/9$ and A_2 with probability $2/9$, in the second she chooses A_3 with probability $1/7$ and A_4 with probability $6/7$, and in the third information set she always chooses A_5 (probability 1) and never A_6 (probability 0). We can show this is a reasonable choice—part of a Nash equilibrium. The probability $p(A_1 A_3 \bullet)$ for reduced strategy $A_1 A_3 \bullet$ would be $p(A_1 A_3 \bullet) = (7/9) \cdot (1/7) = 1/9$. Similarly $p(A_1 A_4 A_5) = (7/9) \cdot (6/7) \cdot 1 = 2/3$, $p(A_1 A_4 A_6) = (7/9) \cdot (6/7) \cdot 0 = 0$, and $p(A_2 \bullet \bullet) = (2/9)$. The numbers sum to 1.

32.1.2 Calculating Behavioral Strategies from Mixed Strategies for a Game Tree with Perfect Recall

Each behavioral strategy gives rise to a mixed strategy. As mentioned, there are more mixed strategies than behavioral ones. So mixed strategies can be distinguished in a finer way than behavioral strategies. Mixed strategies allow coordination among choices in information sets.

Example 1 continued Let me illustrate this by using the VNM POKER(2, 4, 2, 3) example with the extensive form in Figure 31.1. Ann has four pure strategies: RR to raise always, RC to raise with a card of value 1 and check with a 2, CR to check with a card of value 1 and to raise with a 2, and CC to check always. Now consider two mixed strategies: 50% RR, 50% CC and 50% RC, 50% CR. In both Ann passes in 50% of the cases when she has 1 and in 50% of the cases when she has a 2. That would also be Ann's behavioral strategy in both cases. An observer would not be able to distinguish the two. Their difference lies in a different kind of coordination, one that cannot be observed but can be revealed if an observer asks questions. For example, in the game Ann, having raised when playing the first mixed strategy, would answer the question "What would you have done if you had the other card?" like this: "I would also have raised". Ann, having just raised and playing the second mixed strategy, would answer "I would have checked".

We want to convert mixed strategies into equivalent behavioral strategies. This is possible only in games with **perfect recall**, which is the reasonable assumption that a player remembers everything that she observed in the game, including her own moves.

Assume we are given a mixed strategy for Ann. We want to find an equivalent behavioral strategy. Look at an information set I of Ann. We look at its vertices (positions). Ann doesn't know which she faces—after all, this is the point of an information set. Although the moves of the other players and the random moves so far may not be known to her, she recalls all the information sets that she has faced—call this her information sets history—and she recalls the options she has chosen—the options history. Therefore she can exclude all vertices having a different information sets history or a different options history than those she's experienced. We consider only game trees, so both these histories are unique for every vertex. This means that all vertices of the information set are reached by the same sequence of Ann's decisions—as "first choose option A_3 in

information set 2, then choose option A_7 in information set 6, and then choose option A_{13} in information set 5.” In this case the information set history would consist of information sets 2, 5, and 6.

To define probabilities for options in the information set I , we look at all pure strategies that agree with I ’s options history on I ’s information set history. Their probabilities are added to give a number M . Then, for every option A_j of I , we add the probabilities of the pure strategies agreeing with I ’s option history on I ’s information set history, and choosing A_j in information set I . Denote the sum by N_j . Then $p(A_j) = N_j/M$.

Example 1 continued In VNM POKER(2, 4, 2, 3) (Figure 31.1) both of Ann’s information sets have no history. Whether or not they are reached depends on a random move, but not on Ann’s previous moves, since she doesn’t have any. Now consider a mixed strategy of Ann of p_1 RR, p_2 RC p_3 CR, and p_4 CC, where $p_1 + p_2 + p_3 + p_4 = 1$. The corresponding behavioral strategy raises when facing a 1 with probability $p_1 + p_2$, and with probability $p_1 + p_3$ when facing a 2. Now The mixed strategies

- $p_1=30\%$, $p_2=10\%$, $p_3=50\%$, and $p_4=10\%$
- $p_1=20\%$, $p_2=20\%$, $p_3=60\%$, and $p_4=10\%$
- $p_1=10\%$, $p_2=30\%$, $p_3=70\%$, and $p_4=10\%$

all lead to the same behavioral strategy of raising in 40% of the cases with a 1, and in 80% of the cases with a 2. If we translate the behavioral strategy back into mixed strategies, we get $p_1 = p(RR) = 40\% \cdot 80\% = 32\%$, $p_2 = p(RC) = 40\% \cdot 20\% = 8\%$, $p_3 = p(CR) = 60\% \cdot 80\% = 48\%$, and $p_4 = p(CC) = 60\% \cdot 20\% = 12\%$. The mixed strategy also has $p(RR) + p(RC) = 40\%$ and $p(RR) + p(CR) = 80\%$.

Example 2 continued As a second example, in the game in Figure 32.1, let us translate into a behavioral strategy Ann’s mixed strategy $p(A_1 A_3 \bullet) = 0$, $p(A_1 A_4 A_5) = 0.3$, $p(A_1 A_4 A_6) = 0.3$, and $p(A_2 \bullet \bullet) = 0.4$. This is a Nash equilibrium mixed strategy for Ann. Ann’s first information set has no history, the second information set has the first information set with move A_1 as history, and the third information set has both other information sets with moves A_1 (in information set 1) and A_4 (in information set 2) as history.

What would Ann do in the first information set? We consider all pure reduced strategies. We add the three probabilities $p(A_1 A_3 \bullet)$, $p(A_1 A_4 A_5)$, and $p(A_1 A_4 A_6)$ for pure reduced strategies choosing A_1 to get the probability $0 + 0.3 + 0.3 = 0.6$ for A_1 . In the same way, $p(A_2) = p(A_2 \bullet \bullet) = 0.4$.

The second information set can be reached only when Ann chooses A_1 in the first information set, only with pure strategies $A_1 A_3 \bullet$, $A_1 A_4 A_5$, and $A_1 A_4 A_6$, which occur with probability $0 + 0.3 + 0.3 = 0.6$. Then

$$p(A_3) = \frac{p(A_1 A_3 \bullet)}{0.6} = \frac{0}{0.6} = 0$$

and

$$p(A_4) = \frac{p(A_1 A_4 A_5) + p(A_1 A_4 A_6)}{0.6} = \frac{0.3 + 0.3}{0.6} = 1.$$

The third information set can be reached only when Ann plays A_1 and then A_4 , only with pure strategies $A_1 A_4 A_5$ and $A_1 A_4 A_6$, occurring with probability $0.3 + 0.3 = 0.6$. We get

$$p(A_5) = \frac{p(A_1 A_4 A_5)}{0.6} = \frac{0.3}{0.6} = 0.5,$$

and similarly

$$p(A_6) = \frac{p(A_1 A_4 A_6)}{0.6} = \frac{0.3}{0.6} = 0.5.$$

32.1.3 Kuhn's Theorem

We now state Kuhn's Theorem, which summarizes the correspondence between mixed and behavioral strategies.

Theorem [K1953] *In every game with perfect recall, every behavioral strategy is equivalent to one or infinitely many mixed strategies, and every mixed strategy is equivalent to exactly one behavioral strategy.*

Exercises

1. Assume there are the following two behavioral strategies in MYERSON POKER (displayed in Figure 24.8) for Ann and Beth. Ann would always raise when having a red card, and with a blue card she would raise in $1/3$ of the cases. Beth would call in $2/3$ of the cases. What are the expected payoffs for Ann and Beth? Translate the behavioral strategies into mixed strategies for Ann and Beth.
2. In MYERSON POKER, Ann has four pure strategies (raise, raise), (raise, check), (check, raise), and (check, check), where the first entry refers to Ann's red card vertex, and the second to the blue card vertex. Beth has the two pure strategies (call) and (fold). Translate Ann's mixed strategy that uses the four pure strategies with probability $1/4$ into a behavioral strategy.

The next questions refer to LE HER*(4, 2), explained in Chapter 24. The game is played with an eight-card deck, two aces, two kings, two queens, and two jacks. Ann and Beth both get a card at which they look without showing it to their opponent. A third card is put, face down, on the table. If Ann doesn't hold an ace, she can decide whether she wants to exchange cards with Beth. If Beth doesn't hold an ace, she has the opportunity to exchange her card with that lying on the table. The players reveal their cards, and the higher one wins.

3. Let Ann play the behavioral strategy of always exchanging the card if it is a jack, exchanging with probability $1/3$ if it is a queen, and never changing an king. What is the corresponding mixed strategy for Ann?
4. Beth plays the following behavioral strategy: If Ann didn't exchange her card, then Beth always exchanges a jack, exchanges a queen with probability $2/3$, and never exchanges a king. If Ann exchanged, then there are nine situations: Beth will then never exchange if she has a king, and always if she has a jack. If she has a queen, she will not exchange if Ann has a jack or queen, but always if Ann has a king or an ace. Compute the corresponding mixed strategy for Beth.
5. Let $X_J X_Q X_K$ be a pure strategy for Ann, where X_J , X_Q , and X_K refer to the information sets holding jack, queen, or king. The entries are C for exchanging cards with Beth, and N for not exchanging cards. What is an equivalent behavioral strategy for choosing CCN in 30% of the cases, NNN in 30% of the cases, and CNN in the remaining 40% of the cases?
6. Pure strategies for Beth have fifteen entries, since she has fifteen different information sets. We combine Beth's decision into a fifteen-letter word

$$X_J X_Q X_K X_{JJ} X_{JQ} X_{JK} X_{QJ} X_{QQ} X_{QK} X_{KJ} X_{KQ} X_{KK} X_{AJ} X_{AQ} X_{AK},$$

with letters C or N depending on whether she exchanges with the unknown card on the table. The letters X_j refer to cases where Ann didn't exchange cards and Beth has card j , and the letters X_{ij} refer to

cases where Ann exchanged and Ann has card i and Beth card j . Assume Beth chooses CCNCNNC-NNCCNCCC with probability $1/2$, and CCNCNNCCNCCNCNC and CNNCNNCNNCCNNCN with probability $1/4$. Find an equivalent behavioral strategy.

7. Show that Ann's strategy in question 3 is never part of a Nash equilibrium.
8. Show that Beth's strategy in question 4 is never part of a Nash equilibrium.

CHAPTER 33

Example: Multiple-Round Chicken

Prerequisites: Chapters 12, 22, 24, 27, and 32.

In October 1962, US surveillance discovered that the USSR was about to install offensive atomic missiles in Cuba. For one week President John F. Kennedy and his advisors kept the knowledge secret and discussed what to do. It has been reported that in these meetings the Secretary of Defense, Robert McNamara, outlined three options: trying to solve the problem politically, continuing the surveillance and starting a blockade, or undertaking military action against Cuba. A week after the discovery, when it released information about the crisis, the US announced it would undertake a blockade, which it called a “quarantine”. As Russian ships continued to steam toward the blockade, many around the world thought that the US and USSR (and hence all nations) would slip into an atomic war. After several days of blockade, letters between Kennedy and Soviet Leader Nikita Khrushchev, and diplomacy behind the scenes, the conflict was resolved. The Soviets removed their offensive missiles from Cuba.

Although “game” doesn’t seem appropriate for a crisis as serious and threatening as this one, game theory and in particular the game of CHICKEN is often used to model the Cuban missile crisis. In CHICKEN both parties move simultaneously, move only once, and have just two options: to stand firm, the Hawk move, or to give in, the Dove move. If both stand firm, a nuclear war is unavoidable. As far as the payoffs are concerned, at that time, 1962, it was obvious that a nuclear war would be by far the worst possible outcome for both parties. Both giving in does not change things much, but if one gives in and the other remains firm, the firm party has a slight advantage, both for achieving its goal, and raising its stature in public opinion.

Such a simple model can be criticized. Isn’t it a little odd that the players have only one move which they perform simultaneously when the conflict lasts two weeks, with several opportunities every day for both parties to give in? Maybe the conflict should be described as a multiple step game, where, if you stand firm, you have to confirm your decision repeatedly. In the rounds the probability of disaster—atomic war—would increase, until eventually such a war would be unavoidable if none of the parties changed its mind. Why would we introduce random moves in a game like this? Wasn’t the decision to start a war one that could have been made only by the President or the Soviet Premier? In my opinion the conflict teaches us that this is not the case. There were events in the Cuban missile crisis, for instance Soviet boats approaching the blockade line, a US U-2 spy airplane accidentally entering Soviet airspace, or a U-2 spy airplane being shot down over Cuba by a missile ordered by a Soviet commander, that could have triggered nuclear war. One wrong reaction, and Kennedy or Khrushchev would have been left with no choice but all-out war.

33.1 Ordinary Chicken

Here is the conflict modeled as an ordinary chicken game. If both players play Dove, the Russians pull back their missiles from Cuba and the Americans lift their naval blockade, so nothing much happens. The model gives payoffs of 0 for both players. If both players play Hawk, then an atomic war occurs, with payoffs of

−10 for both. It probably wouldn't matter who wins the war. In the mixed cases the Hawk player has an advantage, expressed by a payoff of 1 versus a slight loss, a payoff of −1, for the player playing Dove. The bimatrix is:

	Dove	Hawk
Dove	0, 0	−1, 1
Hawk	1, −1	−10, −10

There are two pure Nash equilibria, the combinations of Dove versus Hawk. There is a mixed Nash equilibrium with 10% Hawk for both, giving both a payoff of −0.1, with war occurring in $10\% \cdot 10\% = 1\%$ of the cases. Since it is the only symmetric Nash equilibrium for the game, it might be the most likely to occur.

We assume an interval scale for the payoffs if we bring in mixed strategies. In the next two sections, where randomness and expected payoffs occur, we will assume an interval scale even for the pure Nash equilibria. In all our models, payoffs of −11 for the atomic war outcome would change the results. But since the results would be changed in the same direction, we'll assume the values of −10.

33.2 Two-round Chicken

We now model the conflict as a multi-step sequential game with imperfect information and randomness, as in Figure 33.1. Why we use a sequential game with imperfect information, or with simultaneous moves in each round, can be explained by the lack of reliable communication channels between the USSR and US leaders. It took hours for a decision to reach the opponent. We assume that if both players play Hawk in the first round, then there is a 25% probability of war after the first round.




Figure 33.1. The Game Tree of two-round CHICKEN

Student Activity Play the game several times with different players. Discuss the model. What would you expect? Would the payoffs for the opponents and the probability of a war increase or decrease when playing Nash equilibria?

We will analyze the game twice, in 33.2.1 using generalized backward induction and the extensive form, and in 33.2.2 by finding pure reduced strategies and the normal form.

33.2.1 Generalized Backward Induction, using the Extensive Form

This game has the ordinary chicken game as a subgame, the part starting with Adam's second information set, which is a single vertex. Therefore we can use the generalization of backward induction discussed in Chapter 22 to attack the problem. The plan is to analyze the subgame first and then put in the values obtained from any Nash equilibria we find as payoffs into Adam's second vertex to find the Nash equilibria for the whole game.

We first concentrate on assigning payoffs to Adam’s second move. Arriving at that situation means that both players have stood firm in the first round, and they had luck—war didn’t start yet. Since at this point we face ordinary one-round chicken, that situation’s three Nash equilibria transform into three cases at Adam’s second vertex. So we have payoffs there of $-1, 1$; or $1, -1$; or $-1/10, -1/10$. The next step is to compute the expected payoffs at the random vertex. In all three cases we take the weighted average of the $-10, -10$ payoffs in case of war and the payoffs at Adam’s second move. We get payoffs of

- $(1/4) \cdot (-10) + (3/4) \cdot (-1)$ versus $(1/4) \cdot (-10) + (3/4) \cdot (1)$, yielding $-3.25, -1.75$,
- $(1/4) \cdot (-10) + (3/4) \cdot (1)$ versus $(1/4) \cdot (-10) + (3/4) \cdot (-1)$, yielding $-1.75, -3.25$,
- $(1/4) \cdot (-10) + (3/4) \cdot (-1/10)$ versus $(1/4) \cdot (-10) + (3/4) \cdot (-1/10)$, yielding $-2.575, -2.575$.

All that remains to do is to consider the first round, with these three possibilities of payoffs for the random move. We get three simultaneous games:

Dove versus Hawk in round 2:

	Dove	Hawk
Dove	0, 0	-1, 1
Hawk	1, -1	-3.25, -1.75

Hawk versus Dove in round 2:

	Dove	Hawk
Dove	0, 0	-1, 1
Hawk	1, -1	-1.75, -3.25

Both mix in round 2:

	Dove	Hawk
Dove	0, 0	-1, 1
Hawk	1, -1	-2.575, -2.575

Let’s analyze the games. In the first game we get the usual two Dove versus Hawk pure Nash equilibria, and one mixed one where Adam stands firm with probability $4/7 = 0.571$ and Bob stands firm with probability $4/13 = 0.308$. The first case corresponds to Adam playing Dove and Bob playing Hawk in the second round. The second game is similar, but with the roles of Adam and Bob reversed. The third game also has these two pure Nash equilibria, and a mixed one where both stand firm with probability $40/103 = 0.388$.

We get a total of nine subgame-perfect Nash equilibria:

- Adam plays Dove in the second round, and Bob plays Hawk in the second round.
 - (1) Adam plays Dove in the first round, and Bob plays Hawk in the first round, DD versus HH.
 - (2) Adam plays Hawk in the first round, and Bob plays Dove in the first round, HD versus DH.
 - (B3) Adam plays Hawk in the first round with probability $4/7 = 0.571$, and Bob plays Hawk in the first round with probability $4/13 = 0.308$. Adam’s and Bob’s expected payoffs are $-4/7$ and $-4/13$, and the probability for war is 0.044.
- Adam plays Hawk in the second round, whereas Bob plays Dove in the second round.
 - (4) Adam plays Dove in the first round, and Bob plays Hawk in the first round, DH versus HD.
 - (5) Adam plays Hawk in the first round, and Bob plays Dove in the first round, HH versus DD.
 - (B2) Adam plays Hawk in the first round with probability $4/13 = 0.308$, and Bob plays Hawk in the first round with probability $4/7 = 0.571$. Adam’s and Bob’s expected payoffs are $-4/7$ and $-4/13$, respectively. The probability for war is 0.044.

- Adam and Bob both play Hawk in the second round with probability $1/10$.
(B5) Adam plays Dove in the first round, and Bob plays Hawk in the first round.
(B4) Adam plays Hawk in the first round, and Bob plays Dove in the first round.
(B1) Both Adam and Bob play Hawk in the first round with probability $40/103 = 0.388$. The expected payoff for both players is $-40/103$, and the probability for war is 0.039 .

There are four pure strategies (1), (2), (4), and (5) and five Nash equilibria with behavioral strategies (B1) to (B5). Strategies (B4) and (B5) may look a little funny. They describe what would happen in the second round, although the second round will not occur.

Interestingly, in the equilibria (B1)–(B3) the payoffs for both players are smaller than in the mixed Nash equilibrium in the one-round model, and the probabilities for war are larger. Thus there seems to be no advantage for the world when moving to the two-step model if (B1), (B2) or (B3) occurs! Maybe a conflict should rather be short and very dangerous than long with increasing danger?

33.2.2 Working with the Normal Form

To practice, let us analyze the game again, this time using its normal form: Both players have two information sets, the options D (Dove) and H (Hawk) in each, and therefore have four pure strategies: DD, DH, HD, and HH. However, if a player plays Dove in the first round, then it doesn't matter what he would have played in the second round. Accordingly the strategies DD and DH are indistinguishable and are subsumed under the reduced strategy D•. The reduced normal form is

	D•	HD	HH
D•	0, 0	−1, 1	−1, 1
HD	1, −1	−5/2, −5/2	−13/4, −7/4
HH	1, −1	−7/4, −13/4	−10, −10

When we solve the game, using any method discussed in Chapter 27, we get four pure Nash equilibria: D• versus HD, D• versus HH, HD versus D•, and HH versus D•. There are three mixed Nash equilibria, and two infinite families of half-mixed Nash equilibria:

- (M1) $63/103$ of D•, $36/103$ of HD, and $4/103$ of HH versus the same mixed strategy,
- (M2) $9/13$ of D• and $4/13$ of HH versus $3/7$ of D• and $4/17$ of HD,
- (M3) $3/7$ of D• and $4/17$ of HD versus $9/13$ of D• and $4/13$ of HH.
- (M4, a) a of HD and $1 - a$ of HH versus D•.
- (M5, a) D• versus a of HD and $1 - a$ of HH.

The Nash equilibria in pure strategies are extreme cases of the families (M4, a) and (M5, a)—they are (M4, 0), (M4, 1), (M5, 0), and (M5, 1). The Nash equilibria of the form (M4, a) and (M5, a) are weak in the sense that a deviation from the Nash equilibrium by one player (the one playing the mix of HD and HH) does not necessarily reduce the payoff for that player: it may keep it constant.

33.2.3 Connections between the two Approaches

The solutions computed in sections 2.1 and 2.2 should be identical. This is obvious for the pure strategies, but how do the five behavioral strategies (B1)–(B5) relate to the three mixed strategies (M1)–(M3)?

Let's translate (B1)–(B5) into mixed strategies, using the formula on the page about behavioral strategies. Then (B1) translates into both players choosing D• with probability $63/103$, HD with probability $(40/103) \cdot (9/10) = 36/103$, and HH with probability $(40/103) \cdot (1/10) = 4/103$. This is mixed strategy (M1). In the same way, the pairs of behavioral strategies (B2) and (B3) translate into the pairs of mixed strategies (M2) and

(M3). (B4) translates into Adam playing 9/10 of HD and 1/10 of HH, and Bob playing the pure strategy D. This is (M4, 9/10). In the same way, behavioral strategy (B5) translates to the mixed strategy (M5, 9/10).

Conversely, let us translate the mixed strategies into behavioral strategies. At each player's first vertex the probability for Hawk should be $P(\text{HD}) + P(\text{HH})$, and the probability for Dove should be $P(\text{D}\bullet)$. At the second vertex, the probability for Hawk should be $P(\text{HH})/(P(\text{HD}) + P(\text{HH}))$, and the probability for Dove should be $P(\text{HD})/(P(\text{HD}) + P(\text{HH}))$. If we translate the five mixed strategies, we get:

- (B1) Both players play Hawk with probability $40/103$ at their first vertex, and with probability $1/10$ at their second vertex.
- (B2) Adam plays Hawk with probability $4/13$ at his first vertex, and with probability 1 (that is, always) at his second vertex. Bob plays Hawk with probability $4/17$ at his first vertex, and with probability 0 (that is, never) at his second vertex.
- (B3) Like (B2), with Adam and Bob interchanged.
- (C1) Adam always plays Hawk at his first vertex, and Hawk with probability $1 - a$ at his second vertex. Bob never plays Hawk on his first vertex, and Hawk with any probability (since we obtain the ratio $0/0$) at his second vertex.
- (C2) Like (C1), with Adam and Bob interchanged.

The cases (C1) and (C2) are subgame-prefect (see Section 24.4) only if they coincide with (B4) and (B5).

33.3 Three-round Chicken

We will model how the probability of war increases as time goes by. We assign $1/4$ and $1/2$ as the probabilities for war after the first and the second rounds. If the third round is reached and both parties still stand firm, war is certain. Figure 33.2 displays the extensive form of the three-round version.




Figure 33.2. The game tree of three-round CHICKEN

Since players have three information sets with two choices in each, each player has $2 \cdot 2 \cdot 2 = 8$ pure strategies. We abbreviate them by three letter words made up of H and D, where the first letter refers to the choice in the first round, and so on.

Some pure strategies can be identified, since they lead to identical outcomes. For instance, the four strategies DDD, DDH, DHD, and DHH can be subsumed under the reduced strategy $D\bullet\bullet$. If a player chickens in the first round, the other two rounds will not occur. Similarly, the pure strategies HDD and HDH are combined under $HD\bullet$. Each player has four reduced strategies, $D\bullet\bullet$, $HD\bullet$, HHD, and HHH. Here is the reduced normal form:

	D●●	HD●	HHD	HHH
D●●	0, 0	−1, 1	−1, 1	−1, 1
HD●	1, −1	−10/4, −10/4	−13/4, −7/4	−13/4, −7/4
HHD	1, −1	−7/4, −13/4	−50/8, −50/8	−53/8, −47/8
HHH	1, −1	−7/4, −13/4	−47/8, −53/8	−10, −10

The Nash equilibria are

1. Adam plays D●●, and Bob plays one of the pure strategies HD●, HHD, HHH, or mixes them. The payoffs for Adam and Bob are −1 and 1. The situations where Adam and Bob reverse their play also form Nash equilibria.
2. Adam plays D●● with probability 3/7, and HD● with probability 4/7. Bob plays either 9/13 D●● and 4/13 HHD, or 9/13 D●● and 4/13 HHH, or a mix between them. The expectations for Adam and Bob are $-4/13 \approx -0.31$ and $-4/7 \approx -0.57$, respectively. The reversed situations are also Nash equilibria.
3. Adam chooses D●● with probability 5/8, HD● with probability 7/24, and HHD with probability 1/12. Bob chooses D●● with probability 18/29, HD● with probability 9/29, and HHH with probability 2/29. Adam and Bob expect $-11/29 \approx -0.379$ respectively $-3/8 = -0.375$. The reversed situations are also Nash equilibria.
4. Both Adam and Bob choose D●● with probability 333/535, HD● with probability 162/535, HHD with probability 36/535, and HHH with probability 4/535. Both expected payoffs are $-202/535 \approx -0.378$.

Let us calculate the behavioral strategies for (4): The players choose Hawk at the first move with probability $(162 + 36 + 4)/535 = 202/535 \approx 0.38$, at the second move with probability $(36 + 4)/202 = 40/202 \approx 0.2$, and at the third move with probability $4/40 = 0.1$. Using the behavioral strategies we can calculate the probability for nuclear war when they are played. The probability that both play Hawk in the first round is $\frac{202}{535} \cdot \frac{202}{535} \approx 0.143$. With probability of about $\frac{1}{4} \cdot 0.143 \approx 0.0364$ nuclear war occurs right after this first round. Both play Hawk in the second round with probability $\frac{3}{4} \cdot 0.143 \cdot \frac{40}{202} \cdot \frac{40}{202} \approx 0.0042$. With probability $\frac{1}{2} \cdot 0.0042 \approx 0.0021$ we get nuclear war after the second round. With probability $\frac{1}{2} \cdot 0.0042 \cdot \frac{1}{10} \cdot \frac{1}{10} \approx 0.000021$ both play Hawk in the third round, also resulting in nuclear war. Therefore the total probability for nuclear war is $0.0356 + 0.0021 + 0.000021 \approx 0.0378 = 3.78\%$.

The pure Nash equilibria in (1) will never create nuclear war. However, they and the mixed Nash equilibria in (2) may be less likely to occur, since one player has a small expected payoff, much smaller than those in (3) and (4). The probabilities for nuclear war are 3.77% in (4); and 3.78% in (3). This is slightly smaller than the probability of 3.9% for nuclear war in the two-round version if both play the symmetric mixed Nash equilibrium.

So, do more rounds imply a safer world? Not necessarily. Moving from the two-round model to the three-round model, we added a decision between the first and second round with nuclear war occurring with probability $p = 1/2$ if both choose Hawk. What if we add a decision earlier, with probability $p = 1/8$? That is, what if the probabilities in rounds 1 and 2 are not 1/4 and 1/2 but 1/8 and 1/4? In the three-round version, in the three most relevant mixed Nash equilibria, both players have expected payoffs of about −0.73, and the probabilities for nuclear war are about 6.9%.

Student Activity Discuss whether there is any relevance to real life in the models. Doesn't the analysis imply that there is an increased probability for war in the two-round model? Stake out your position between "It is true—a mathematician has computed it" and "Mathematics and game theory cannot be applied in real-world situations like this."

Exercises

1. In the three-round model compute the probabilities for war in the Nash equilibria of type (2) (Adam plays D●● with probability $3/7$, and HD● with probability $4/7$. Bob plays either $9/13$ D●● and $4/13$ HHD, or $9/13$ D●● and $4/13$ HHH, or any mix of the two pure strategies).
2. One may question the assumption that the players make simultaneous moves in rounds. Couldn't the crisis be described more accurately as a sequential game with randomness and perfect information, where the two players alternate? When a player chickens, he loses 1 and the other one wins 1. When a player yields, war is caused by a probability that increases over time (like $1/5, 2/5, 3/5, 4/5, 1$), or the other player is about to move. Model this with Adam starting to move, then Bob, Adam, Bob, and Adam moving, solve it, and explain why this model is not very useful.

Project 56

Suppose that the outcome of atomic war has a payoff of $-k$ for both players. Analyze the two-round version, and if possible the three-round version, for the parameter k . How do the probable outcomes, the expected payoffs for the players, and the probabilities for atomic war change if k increases?

Project 57

The Cuba crisis situation is not symmetric, but the games discussed are. Modify the payoffs for win and loss for the not-atomic war outcomes such that the game is no longer symmetric. Justify your choice. Then analyze the two-round version, and if possible the three-round version.

Project 58

Discuss the three-round version where the probabilities for nuclear war after both choose Hawk are $1/3$ and $2/3$ after the first and second round. Compare the two three-round versions.

CHAPTER 34

Example: DMA Soccer II

Soccer is not mathematics. You can't calculate everything.

— Karl Heinz Rummenigge, Chair of the Board of Bayern Munich,
November 2007, [*Süddeutsche Zeitung*]

I hope I know the basic math of soccer and I try to apply that.

— Ottmar Hitzfeld, Coach of Bayern Munich in November 2007.
Hitzfeld has a degree in mathematics. [*Süddeutsche Zeitung*]

Prerequisites: Chapters 12, 14, 22, 24, and 27.

Coaches of team sports like soccer supervise the training of their teams, and they play in the sense that they select the players and the system their team will use. Besides trying to motivate their teams, coaches also have the important task of reacting to what they see by substituting players or changing the system.

In a Champion's League soccer game between Bolton Wanderers and Bayern Munich in November 2007, Munich was leading 2-1 when their coach Ottmar Hitzfeld substituted for two key players on his team, Frank Ribery and Bastian Schweinsteiger, taking them out. After that, the Bolton Wanderers tied the score at 2-2, and the game ended there. The substitution occasioned Rummenigge's indirect criticism of Hitzfeld's move, which prompted Hitzfeld's response.

In this section we will extend the simple static simultaneous soccer model discussed in Chapter 14 into a game with three parts, where in each part the coaches move simultaneously. We will see that there are too many pure strategies for the normal form approach, so we will try a variant of backward induction on subtrees.

The game is a variant of static nine-round DMA soccer. First each coach selects a starting distribution for her team. What differs from static DMA soccer is that after the third and sixth rounds the distributions can be changed. One player from each field can be moved to an adjacent field. The changes are made simultaneously by both coaches.

Class Activity Play ten games against your neighbor. Use the applet [DMAdyn](#).

34.1 Multi-round Simultaneous Games

The most obvious feature of the game is its structure of simultaneous moves performed in rounds. But since for DMA soccer we have used the notion of rounds already, let's call them thirds.

The two coaches, Ann and Beth, move three times. Their first, second, and third moves are performed simultaneously.

- First third: At the very beginning, Ann and Beth choose their start distribution.

- Second third: After three rounds, they adjust their distribution simultaneously. Players stay where they are or one player moves from a field to an adjacent field.
- Last third: After six rounds, both adjust their distribution simultaneously, using the same rule.

Between these moves are six random moves where the ball moves left or right. The random moves in a third can be combined into one random move whose outcome is up to 3 goals.

For the extensive form we avoid simultaneous moves by assuming that Ann moves always first, but Beth doesn't know Ann's decision until after Beth has moved. That implies that Ann has single-position information sets but Beth's information sets contain several positions. Each of Ann's single-position information sets is the start of a subgame. So we have subgames starting at Ann's positions at the beginning of the last third, larger subgames starting at Ann's positions at the beginning of the second third, and the full game starting at Ann's first third start position.

Although the game does not have perfect information, and therefore backward induction can not be performed, we can use the same idea, analyzing the game recursively starting at the branch subgames. First we assign expected payoffs for Ann's single-position information sets at the beginning of the last third. The analysis also includes move recommendations for Ann and Beth. Then we assign expected payoffs to Ann's single-position information sets at the beginning of the second third. To be able to do this, we need the values of her expected payoffs for the third move positions. After we have done this for the second and last thirds, we can assign expected payoffs for the start position. Unlike backward induction, the payoffs from the backward analysis need not be unique: it is possible to get different sets of values at some positions, which complicates matters. But in our game, this problem will not occur.

34.2 Information Sets and Moves

The distributions DX, MX, AX for the number of players in defense, midfield, and attack, are encoded as (DX, MX) , since AX follows from $DX + MX + AX = 5$. The (undirected) graph in Figure 34.1 shows the changes in distributions at each move. Two distributions are connected by an edge if they can be transformed into each other. Others, like $(1, 2)$ and $(2, 2)$ would require two moves. You would have to move a player from attack to defense, but these fields are not adjacent, so you would have to move one player from attack to midfield and later one from midfield to defense. The allowable moves are symmetric. If you can transform distribution 1 into distribution 2, then you can also transform distribution 2 into distribution 1 by doing the reverse.




Figure 34.1. The graph of possible changes

At their first move, both players have only one information set. At the second and third moves, the information sets are in principle defined by a combination of the current score and the two distributions. In our payoff system, where all that counts is winning, drawing, or losing, a current score of 1-2 would lead to the same move as a score of 0-1. In each case your team needs two more goals than the opponent to win, and one goal for a draw. So all that matters about the current score is the goal difference, so we merge information sets with the same goal difference.

After we do the merge, we can show that at the second move Ann and Beth each have $7 \cdot 21 \cdot 21 = 3087$ information sets. Ann and Beth may have chosen any of the 21 start distributions, and the goal difference after three rounds might be any of the seven values 3, 2, 1, 0, -1, -2, -3 (viewed as in favor for Ann). At their third move, Ann and Beth each have $13 \cdot 21 \cdot 21 = 5733$ information sets, since the possible goal difference after six rounds could be any of the thirteen values 6, 5, ..., -5, -6. Therefore analyzing, or even formally describing, the game takes a huge amount of effort and space. Since at the first move the players have 21 options, and at the second and third move they have at least two options (in most cases they have five options: a defense player could move to midfield, a midfield player could move to either defense or attack, an attack player could move to midfield, or there is no change), each player has at least $21 \cdot 2^{3087} \cdot 2^{5733}$ pure strategies, a very huge number with more than 2600 digits! This is with the conservative estimate that has two moves in each information set. The number of reduced pure strategies would be much smaller, but using the normal form to analyze the game is not an option.

Therefore we stay with the extensive form, which is also huge, with thousands of vertices. As the methods used for normal form are out of our reach, the recursive method for multi-round simultaneous games is our only option. We will not solve the game completely, but we will see how to do it with more time.

34.3 The Optimal Third Move in Selected Cases

Assume that six rounds have been played and that the coaches are about to make their third move, so we are beginning a last-third subgame. What distributions can be achieved depends on the team's present distribution. The decision on where to move depends on two factors. First, the other team's current distribution. Second, the current score, or, to be more precise, the current goal difference: a coach whose team is behind would act differently (probably strengthening midfield or attack) from a coach whose team is leading. If a team is behind, its coach would likely opt for a higher probability of scoring, even if it carries with it a higher probability for goals by the other team.

Each of the coaches has up to five options for the third move—keeping the current distribution, or changing it to one of the up to four adjacent distributions as displayed in the graph. For each pair of third moves chosen, the probabilities for the ten possible goal patterns achieved in the last part — 3-0, 2-0, 1-0, 0-0, 2-1, 1-1, 0-1, 1-2, 0-2, 0-3—can be computed in the same way as in DMA soccer I. The outcomes of the game are obtained from the current score by adding the ten goal patterns. The convention is to give a payoff of 3 points for a win, 1 point for a draw, and 0 points for a loss. Accordingly, it is not too hard to compute the expected number of points that follows from combinations of team A's and team B's distributions.

34.3.1 A Detailed Example: (2, 2) versus (3, 1)

Assume that during the middle third team A played distribution (2, 2) and team B distribution (3, 1). Then Ann has five options. She can move to one of the distributions (1, 3) or (2, 1) or (2, 3), or (3, 1), or stay at distribution (2, 2). Beth has the options of moving to (2, 2), (3, 0), (3, 2), (4, 0), or keeping the distribution (3, 1). The cases where one team leads by four or more goals are easy to analyze—there is nothing the team behind can do—so we look only at the goal differences where the outcome is still not determined. There are seven of them.

	(2, 2)	(3, 0)	(3, 1)	(3, 2)	(4, 0)
(1, 3)	0.48, 2.28	0.89, 1.69	0.52, 2.21	0.51, 2.17	0.70, 1.92
(2, 1)	0.37, 2.42	1.49, 1.08	0.51, 2.23	0.45, 2.25	1.22, 1.33
(2, 2)	0.42, 2.35	0.89, 1.69	0.47, 2.27	0.42, 2.30	0.70, 1.92
(2, 3)	0, 3	0, 3	0, 3	0, 3	0, 3
(3, 1)	0.26, 2.56	0.89, 1.69	0.33, 2.46	0.27, 2.52	0.70, 1.92

Table 34.1. Payoff bimatrix for distributions of (2, 2) versus (3, 1) used in the middle third, provided Ann's team is one goal behind

Ann is one Goal Behind

If Ann's team is one goal behind, then the expected payoff bimatrix, with entries rounded, is in Table 34.1.

Let me explain how the entries are obtained for one cell, the upper middle cell where Ann moves to distribution (1, 3) and Beth keeps her distribution (3, 1). As explained in DMA Soccer I, in each of the remaining three rounds the probability for a goal for Ann's team is $p_A = (3/4) \cdot (1/4) = 3/16$ (Ann has a midfield of 3 and Beth a midfield of 1, whereas Ann has an attack of 1 and Beth a defense of 3), the probability for a goal for Beth's team is $p_B = (1/4) \cdot (1/2) = 1/8$, but the most likely eventuality is no goal with a probability of $p_0 = 1 - 3/16 - 1/8 = 11/16$. For the last three rounds combined, ten goal patterns are possible, whose probabilities are shown in Table 34.2.

Goals for Ann/Beth	0	1	2	3
0	$p_0^3 \approx 0.325$	$3p_B \cdot p_0^2 \approx 0.177$	$3p_B^2 \cdot p_0 \approx 0.032$	$p_B^3 \approx 0.002$
1	$3p_A p_0^2 \approx 0.266$	$6p_A \cdot p_B p_0 \approx 0.097$	$3p_A \cdot p_B^2 \approx 0.009$	
2	$3p_A^2 p_0 \approx 0.073$	$3p_A^2 p_B \approx 0.013$		
3	$p_A^3 \approx 0.007$			

Table 34.2. Probabilities for the number of goals in three rounds for distribution (1, 3) versus distribution (3, 1)

The values sum to 1.

Therefore, the probability that in the last three rounds the goal difference will

- increase by 3 in favor of Ann's team is $p_A^3 \approx 0.007$
- increase by 2 in favor of Ann's team is $3p_A^2 p_0 \approx 0.073$
- increase by 1 in favor of Ann's team is $3p_A p_0^2 + 3p_A^2 p_B \approx 0.266 + 0.013 = 0.279$
- not change is $p_0^3 + 6p_A p_B p_0 \approx 0.325 + 0.097 = 0.422$
- increase by 1 in favor of Beth's team is $3p_B p_0^2 + 3p_A p_B^2 \approx 0.177 + 0.009 = 0.186$
- increase by 2 in favor of Beth's team is $3p_B^2 p_0 \approx 0.032$
- increase by 3 in favor of Beth's team is $p_B^3 \approx 0.002$.

Therefore, if distributions (1, 3) versus (3, 1) are chosen for the last third, if Ann's team is one goal behind, she will have won at the end with probability $0.007 + 0.073 = 0.079$; there will be a draw with probability 0.279, and Beth's team will have won with probability $0.422 + 0.186 + 0.032 + 0.002 = 0.642$. The expected number of points for Ann's team is $3 \cdot 0.079 + 1 \cdot 0.279 = 0.516 \approx 0.52$. The expected number of points for Beth's team is $3 \cdot 0.642 + 1 \cdot 0.279 = 2.205 \approx 2.21$. These are the values in the first row and third column in Table 34.1.

This analysis is done automatically in the Excel sheet [DMA3.xlsx](#). Go to the “after 2 thirds” sheet. In the black cells, fill in the current distributions for Ann and Beth and the current goal difference. You can change the number of points for win, draw, and loss in cells L4, M4, and N4. After having changed the values, the expected points bimatrix is displayed. Confirm the bimatrix in Table 34.1 in the Excel sheet.

Let us look at the payoff bimatrix in Table 34.1 again. Not too surprisingly, Ann’s options (2, 3) and (3, 1) are strictly dominated by option (2, 1). Since Ann’s team is one goal behind, she should put more emphasis into attack. Moving one player from midfield to attack is better for Ann than moving one player from attack to midfield, or moving one player from midfield to defense. If we delete these last two rows, then Beth’s move (2, 2), moving one player from defense to midfield, strictly dominates her other moves. When deleting these options, Ann’s move (1, 3) strictly dominates the other moves. Therefore the IESD process results in only one Nash equilibrium, (1, 3) versus (2, 2). Instead of strengthening her attack, Ann will move one defense player to the midfield and Beth will do the same, moving one defense player to midfield, even though her team is leading. This could be interpreted as Beth’s anticipating and trying to counter Ann’s strengthening of the midfield.

Other Goal Differences

The expected numbers of points change if the goal difference after six rounds is different, and the recommendations for the third move might also change. For the calculations of the bimatrix however, we can reuse a lot of work we have done for when Ann is one goal behind. We get the same set of moves, and for each pair of moves the probabilities p_A , p_B , and p_0 are the same, as are the probabilities of changes of the goal difference by 3, 2, 1, 0, −1, −2, or −3. In the example of Ann’s move of (1, 3) versus Beth’s move of (3, 1), these numbers are again 0.007, 0.073, 0.279, 0.422, 0.186, 0.032, and 0.002. What differs is that the probabilities are added differently. If Ann is two goals behind, and they continue with moves (1, 3) versus (3, 1), then Ann will win the game with probability 0.007, draw with probability 0.073, and lose with probability $0.279 + 0.422 + 0.186 + 0.032 + 0.002$, so Ann’s expected number of points is $3 \cdot 0.007 + 1 \cdot 0.073$, whereas Beth’s expected payoff is $3 \cdot (0.279 + 0.422 + 0.186 + 0.032 + 0.002) + 1 \cdot 0.073$.

We can use the Excel sheet to calculate the 5×5 bimatrices for the seven cases. Analyzing them, we obtain the following move recommendations and expected number of points for Ann and Beth.

goal difference	Ann’s move and expectation	Beth’s move and expectation
−3	(1, 3): 0.003	(3, 2): 2.993
−2	(1, 3): 0.068	(3, 2): 2.875
−1	(1, 3): 0.480	(2, 2): 2.280
0	(1, 3): 1.320	(2, 2): 1.320
1	(2, 2): 2.347	(2, 2): 0.417
2	(2, 3): 2.900	(2, 2): 0.053
3	(2, 3): 2.995	(2, 2): 0.002

Table 34.3. Results for (2, 2) versus (3, 1) in the last third for different goal differences

If team A is behind or if the score is tied, team A should move one player from defense to midfield and proceed with distribution (1, 3) in the last third. If team A is leading by one goal, then team A should not change. If team A’s advantage is larger, team A should move one player from attack to midfield. Team B should almost always move one player from the strong defense to the weak midfield, except when team B leads by 2 or 3 goals, in which case team B should move its one attacking player to midfield. If a team has an advantage of 4 or more goals, the result of the game is decided.

34.3.2 A Second Example: (1, 3) versus (2, 2)

Here are the results for different goal differences: Ann should keep her distribution in case the current score is tied. If Ann is behind, she would move one player from midfield to attack, with a little mixing between doing this and not changing the distribution if she is one goal behind. If she leads, she moves one player from midfield to defense. Beth should almost never keep her current distribution (with the exception of a little mixing if she leads by one goal). If Beth leads by two goals or more, she should move her one attacking player to midfield, leaving her without attack. Otherwise she would move one player from defense to midfield.

goal difference	Ann's move and expectation	Beth's move and expectation
−3	(1, 2): 0.008	(2, 3): 2.984
−2	(1, 2): 0.120	(2, 3): 2.784
−1	70% of (1, 2) and 30% of (1, 3): 0.551	85% of (1, 3) and 15% of (2, 2): 2.227
0	(1, 3): 1.344	(1, 3): 1.344
1	(2, 2): 2.280	(1, 3): 0.480
2	(2, 2): 2.832	(1, 3): 0.096
3	(2, 2): 2.984	(1, 3): 0.008

Table 34.4. Results for (1, 3) versus (2, 2) in the last third for different goal differences

The expected values are shown in Figure 34.2. If we compare corresponding values, it is clear that Ann has a slight advantage.




Figure 34.2. Ann's and Beth's expectations in the last third for different goal differences, for (1, 3) versus (2, 2)

34.4 The Optimal Second Move for Seven Positions: (1, 3) versus (2, 2) and any Goal Difference

Assume Ann's team started with distribution (1, 3) and Beth's team with distribution (2, 2). How should both teams proceed, depending on the goal difference after the first third, which may be any number between −3 and 3: Should they change their distributions for the second third, and, if so, how? How many points can Ann and Beth expect now?

The positions can be analyzed in the same way as the positions in the previous section provided we know the expectations of the players in a few third move positions. To be more precise, we should know the expectations for the $5 \cdot 5 \cdot 13 = 325$ positions where Ann's team plays any of the distributions (0, 4), (1, 2), (1, 3), (1, 4), (2, 2) that can be achieved from (1, 3). Beth's team plays any of the positions (1, 3), (2, 1), (2, 2),

(2, 3), (3, 1) that can be achieved from (2, 2) and the goal difference is an integer from -6 to 6 . Fortunately, this is only a fraction of the total number of $21 \cdot 21 \cdot 13 = 5733$ of third move positions. But third move positions with goal difference $-6, -5, -4, 4, 5$, or 6 are easy to analyze—they are lost (or won) no matter what Ann does. So only $5 \cdot 5 \cdot 7 = 175$ positions remain to be analyzed. We have analyzed $2 \cdot 7 = 21$ of them, so the additional task is doable, using the Excel sheet. The data is collected on the tab with name “after 1 third” on the Excel sheet.

We also need to calculate the probabilities $p_A = \frac{MA}{MA+MB} \cdot \frac{AA}{AA+DB}$, $p_B = \frac{MB}{MA+MB} \cdot \frac{AB}{AB+DA}$, and $p_0 = 1 - p_A - p_B$ for the 25 combinations of distributions in the second third. This means we have 25 tables for the probabilities for the number of goals in three rounds, each similar to the table above. We substitute values of p_A , p_B , and p_0 . Therefore, during the second third, the goal difference will

- increase by 3 with probability p_A^3
- increase by 2 with probability $3 \cdot p_A^2 \cdot p_0$
- increase by 1 with probability $3 \cdot p_A \cdot p_0^2 + 3 \cdot p_A^2 \cdot p_B$, (1-0 or 2-1)
- not change with probability $p_0^3 + 6 \cdot p_A \cdot p_B \cdot p_0$ (0-0 or 1-1)
- decrease by 1 with probability $3 \cdot p_B \cdot p_0^2 + 3 \cdot p_A \cdot p_B^2$, (0-1 or 1-2)
- decrease by 2 with probability $3 \cdot p_B^2 \cdot p_0$
- decrease by 3 with probability p_B^3 .

Thus if we select a pair of distributions in the second third, seven positions are possible at move 3, and their probabilities are as described, where p_A , p_B , and p_0 depend on the distributions chosen. Then we get the expected payoffs by multiplying the expected payoffs for the seven positions by the seven probabilities, and adding these seven products. This is done in the “after 1 third” tab in [DMA3.xlsx](#). Only the calculation of the probabilities and the reference to the seven positions is done automatically there— if one would try a different case, one would have to manually insert the data in the 13 tables below the main table, which is a tremendous amount of work—solving $5 \cdot 5 \cdot 7$ third move cases!

In our example the matrices for the seven second move positions can be created on this tab, by changing the “goal advantage for Ann” entry, which is the only cell value you can change on this tab. Analyzing the seven games, we see that all but one have a pure Nash equilibrium. The mixed Nash equilibrium for the remaining case is done using Brown’s fictitious play. We get the best distributions and expected values in Table 34.5.

Figure 34.3 shows the expected values for the various goal differences. Since we are in an earlier stage of the game, the advantage for the leading team is not as large for goal difference 3, 2, or 1 as in the third move case.

goal difference	Ann’s move and expectation	Beth’s move and expectation
−3	(1, 2): 0.117	(2, 2): 2.840
−2	(1, 2): 0.338	(1, 3): 2.559
−1	40% of (1, 2) and 60% of (1, 3): 0.759	60% of (1, 3) and 40% of (2, 2): 2.057
0	(1, 3): 1.390	(1, 3): 1.390
1	(2, 2): 2.066	(1, 3): 0.737
2	(2, 2): 2.593	(1, 3): 0.292
3	(2, 2): 2.873	(1, 3): 0.082

Table 34.5. Results for (1, 3) versus (2, 2) in the second third for different goal differences




Figure 34.3. Ann’s and Beth’s expectations in the second third for different goal differences, for (1, 3) versus (2, 2)

Most of the move recommendations seem to be what we would expect. Moreover, they are almost identical to those we arrived at in the analysis of the third move above, for the same distributions in Section 3.2, with two exceptions. At the end of the second third, when Beth’s team was leading by at least two goals, our recommendation was to shift the only attack player into the midfield. This is not appropriate now. Maybe it is too early in the game to do that. Instead, Beth now keeps her distribution if she is leading by three goals, and interestingly, gets more offensive and orders one player to move from defense to midfield if she is leading by two goals. She mixes the two moves if she leads by one goal.

34.5 Couldn't We Analyze the Whole Game?

Yes, but it would be *very* tedious. We analyzed 7 of the $21 \cdot 21 \cdot 7 = 3087$ second move positions. For this we had to analyze 325 of the $21 \cdot 21 \cdot 13 = 5733$ third move positions. To analyze the start position, we would have to analyze all these positions. This task might be a little too big for Excel sheets. Are any of you programmers?

We know one entry of the 21×21 matrix for the first move—the entry where Ann plays (1, 3) and Beth (2, 2). Using the same procedures we discussed in the previous section, we can conclude that Ann’s expected value is 1.414 and Beth’s 1.400, a slight advantage for Ann playing (1, 3). Compare this to the results from Chapter 14 that (1, 3) and (2, 2) formed a Nash equilibrium for goal difference, but that the only Nash equilibrium in the static three rounds DMA game was (1, 3) versus (1, 3) (provided a win yields 3 points and a draw 1).

34.6 How Good a Model is it?

Models never represent reality exactly. Usually we have to simplify, as we did in this chapter’s example. Soccer is far too complicated to be modeled because the outcome of a real soccer game depends on more than the distributions of the teams on defense, midfield, and attack. Exchanges can be done at any time. Still, game theorists hope that a model’s recommendations can give information about real situations. We may be able to derive from a valid model qualitative recommendations, like “When behind, shift forces from defense to midfield”. Our model seems to do this much.

CHAPTER 35

Example: Sequential Quiz Show II

A glimpse into cooperative game theory

Prerequisites: Chapters 16 and 17.

This chapter provides a glimpse into cooperative game theory, which is otherwise not covered. We will investigate what happens in $\text{SEQUENTIAL QUIZ SHOW}(n, m)$ if two players cooperate, and share evenly, or according to a formula, a win or loss. In Section 35.1 we look at fixed coalitions. In Section 35.2 we ask which coalitions are most likely to form provided the players in one have to share a win or loss evenly. In Section 35.3 we drop the requirement of having to share evenly. In Section 35.4 we investigate what would be a fair share of the win if all three players work together.

Here is the game description again:

SEQUENTIAL QUIZ SHOW(n, m) Three players, Ann, Beth, and Cindy, are facing a difficult multiple choice question with five options. Starting with Ann and continuing cyclically, a player can either try to give an answer or wait. If the player tries an answer and the answer is correct, the player gets $\$n$. If it is incorrect, the player has to pay $\$m$ and is out of the game. If the player waits, the quizmaster reveals a wrong answer (decreasing the number of options by one), and the next player moves.

35.1 Fixed Coalitions

35.1.1 Ann and Cindy Form a Coalition

Assume the candidates have been selected for next day's show. They know that Ann will start, and Beth will move second. Contestants win \$10 for a correct answer and lose \$4 for a wrong one. So we are playing $\text{SEQUENTIAL QUIZ SHOW}(10, 4)$. As the players meet in the hotel lobby the day before the show, it turns out that Ann and Cindy know each other and are good friends, such good friends that they don't care which one of them wins. They promise they will later share evenly a win or loss. They have plenty of time to coordinate their strategies. Does cooperation affect the game? Will the players play differently? Should Beth expect less?

We are not just talking about a change of strategies—the game has totally changed. We are no longer playing the three-player game $\text{SEQUENTIAL QUIZ SHOW}(10, 4)$ but a two-player variant with the coalition as one player and Beth the other. The payoff for the coalition is the sum of Ann's and Cindy's payoffs. We assume that Beth knows about the coalition — otherwise we would have to enter the tricky field of games of incomplete information.




Figure 35.1. Solution of the (10, 4) case if Ann and Cindy form a coalition

Figure 35.1 shows the extensive form with backward induction analysis, where the payoff of the coalition is in red. The only difference from the strategies in the ordinary game is that Cindy waits when it is her turn. This increases the payoff for the coalition from $2\frac{2}{3}$ to 3. It can be seen that Beth profits more from Cindy not trying an answer—Beth’s expected payoff increases from $3\frac{1}{3}$ to 5. So Beth would probably not mind that Ann and Cindy have a coalition. Or would she try to form a coalition with Ann or Cindy instead?

35.1.2 Ann and Beth Form a Coalition

The analysis of this game, a 2-player game of the coalition of Ann and Beth versus Cindy, is left as an exercise, but let us reveal the result. The players of the team play slightly differently from how they would play if they were not a team. Ann and Beth wait, but after Cindy has tried an answer and failed, Ann can wait again and let Beth give the correct answer. In this way, the expected total payoff of the team goes up to $6\frac{2}{3}$, whereas Cindy’s expected payoff stays at $\frac{2}{3}$.

35.1.3 Beth and Cindy Form a Coalition

Ann and Beth will not try an answer. Cindy can also afford not to try an answer, since she knows that Ann will try an answer next but will fail with probability $\frac{1}{2}$, in which case Beth will win for the team. The expected payoff for the team is 5, one unit larger than the sum of Beth’s and Cindy’s payoffs in the non-team case. Even Ann profits from the changed strategy, because her payoff is 3, one larger than in the ordinary case.

35.2 Which Coalition Will Form?

35.2.1 Fixed 50:50 Split

Consider a different variant. Two players may form a coalition and discuss their moves before the game starts provided they announce the coalition before the game starts, and they split the win or loss evenly, 50:50. The two players also apprise the quiz show administration of the distribution of prize money.

For this and the following sections, we use the following terminology: Let $v(A)$, $v(B)$, and $v(C)$ be the expected payoffs for Ann, Beth, and Cindy if the game is played without coalition and let $w(AB)$ be the expected total payoff of the coalition of Ann and Beth, whereas $w(C)$ means the expected payoff for Cindy provided she faces a coalition of Ann and Beth. The values $w(AC)$, $w(B)$, $w(BC)$, and $w(A)$ are defined accordingly.

$n = 10, m = 4$

Suppose $n = 10$ and $m = 4$. We get $v(A) = 2$, $v(B) = \frac{10}{3} \approx 3.33$, $v(C) = \frac{2}{3} \approx 0.67$, and $w(AC) = 3$, $w(B) = 5$, $w(AB) = \frac{20}{3} \approx 6.66$, $w(C) = \frac{2}{3} \approx 0.67$, $w(BC) = 5$, $w(A) = 3$, as seen in the previous section.

The following table shows the expected values for each player after the coalitions have divided the 50:50 payoff.

	no team	Ann and Cindy	Ann and Beth	Beth and Cindy
Ann	2	1.5	3.33	3
Beth	3.33	5	3.33	2.5
Cindy	0.67	1.5	0.67	2.5

Thus if the money is split 50:50, Ann will not agree to an Ann and Cindy coalition, and Beth will not agree to a Beth and Cindy coalition. Thus the only possible coalition is Ann and Beth, although Beth wouldn’t care whether it formed or not.

$n = 12, m = 4$

When $n = 12$, $m = 4$, we get $v(A) = \frac{8}{3} \approx 2.67$, $v(B) = 4$, $v(C) = \frac{4}{3} \approx 1.33$, $w(AC) = 4$, $w(B) = 3.2$, $w(AB) = 8$, $w(C) = \frac{4}{3} \approx 1.33$, $w(BC) = 6$, and $w(A) = 4$.

The next table shows the expected payoffs for each player after the coalitions have divided the payoff 50:50.

	no team	Ann and Cindy	Ann and Beth	Beth and Cindy
Ann	2.67	2	4	4
Beth	4	3.2	4	3
Cindy	1.33	2	1.33	3

The only coalition that could form would be Ann and Beth.

35.3 Another Variant: Split can be Negotiated

In this variant two players may form a coalition and agree on their moves and on how to distribute a win or loss, provided they announce their plans before the game starts. They promise, and the quiz show administration knows how they will split their win or loss.

Should a coalition form? If so, which one? And which split should it use? To investigate these questions, let us concentrate on the negotiation process, which takes place before the quiz show game is played. It makes sense to make a few assumptions about it, and also to formalize it. So assume the three players are sitting around a table and talking about coalitions. Two players can tentatively agree on a coalition and a split ratio. However, at any time only one tentative coalition may exist, and no three player coalition is possible. Two players may agree to a tentative coalition only if it would give them larger expected payoffs. The third player obviously does not have to agree to a tentative coalition. If two players form a tentative coalition that includes a split ratio, the ratio cannot be changed, since for one player change would mean decreasing her payoff. The tentative coalition can terminate in two ways. One player involved announces that she is no longer interested in the coalition. We assume that she would do this only if her payoff without coalition would be higher than her expectations in the coalition. Or, the coalition could terminate if one of the players involved wants to form a tentative coalition with the third player, who agrees to it. The new tentative coalition must increase the payoffs of both of its members.

We say a coalition is **stable** if neither of the two ways of terminating it occurs, because the expected payoffs in the no-coalition case would not be higher for both players belonging to the coalition, and because any new coalition that would involve the third player and give her a higher expected payoff, would necessarily decrease the expected payoff for the other player in the coalition.

The negotiation phase ends if no movement has been achieved for some time, or if the two coalition players firmly believe that the current coalition gives the split they desire. Then contracts are signed and the game is played. We assume that a final coalition is stable.

We express splits of win or loss not as percentages but rather in terms of the expected values. So, if for instance the Ann and Beth coalition would expect a joint payoff of 8, then 5:3 is a possible split, meaning that Ann gets $\frac{5}{8}$ of the win or loss and Beth gets $\frac{3}{8}$.

If Ann does not agree to a coalition, then she expects $w(A)$ or $v(A)$, depending on whether Beth and Cindy form a coalition or not. Therefore we may assume that Ann would not agree to a coalition and split where her expected value is less than the smaller of $v(A)$ and $w(A)$. Let us look at two cases:

$n = 12, m = 4$

Student Activity Form groups of three and simulate the negotiations for the values for $n = 12, m = 4$. What is the result of your negotiations?

We use the results from Section 35.2.1: $v(A) = \frac{8}{3} \approx 2.67, v(B) = 4, v(C) = \frac{4}{3} \approx 1.33, w(AC) = 4, w(B) = 3.2, w(AB) = 8, w(C) = \frac{4}{3} \approx 1.33, w(BC) = 6$, and $w(A) = 4$. Possible tentative coalitions are Ann and Beth with splits from 4.8:3.2 to 2.67:5.33, and Ann and Cindy with the split of 2.67:1.33, and Beth and Cindy with splits from 3.2:2.8 to 4.67:1.33. The Ann and Cindy coalition is not stable, since almost every Ann and Beth coalition is better for both Ann and Beth. However, the Ann and Beth coalitions with splits

from 4.8:3.2 to 3.33:4.67 are also not stable, since some Beth and Cindy coalitions would be better for both Beth and Cindy. For example, for a 4:4 split Cindy expects $w(C) = 1.33$. Now Beth and Cindy could profit from a coalition with a 4.3:1.7 split. The Beth and Cindy coalitions with splits from 3.2:2.8 to 4:2 are also not stable—Beth could terminate the coalition and expect a payoff of 4 in the no-coalition case. Beth does not have to fear the Ann and Cindy coalition, since it is not stable, as we saw above.

Thus the only stable coalitions are Ann and Beth coalitions with splits from 3.33:4.67 to 2.67:5.33, and Beth and Cindy coalitions with splits from 4:2 to 4.67:1.33. These Beth and Cindy coalitions are stable because, though Beth would be interested in switching to a Ann and Beth coalition, Ann would not. These Ann and Beth coalitions are stable since, while Ann would be interested in Beth's forming a coalition with Cindy, Ann cannot initiate one, and Beth is not interested in coalitions with lower expectations. One might think that Ann would terminate the coalition, decreasing her payoff to 2.67, but hoping that Beth and Cindy will form a coalition, assuring Ann a payoff of 4. Such a move is not considered in the negotiation process above—it is too risky.

$n = 10, m = 4$

Student Activity Form groups of three and simulate the negotiations for $n = 10, m = 4$. What is the result of your negotiations?

This case is more confusing. No stable coalition exists. See the project in the exercises.

35.4 The Grand Coalition

Wouldn't it make the most sense for all three players to form a coalition? Ann, Beth, Cindy, and Ann again would wait, until Beth has only one answer left, which will be correct. The payoff for the coalition would be n . We would need enforceable contracts before the game starts to ensure that Beth doesn't take the money and run but shares it with Ann and Cindy according to a scheme they would have written into this contract.

35.4.1 The Core

So far, so good. How would the three players share the payoff of n ? Ann, Beth, and Cindy each should get at least as much as they would get when playing alone, otherwise they may not cooperate. Moreover, for the same reason each pair should get at least as much as they would get when forming two-player coalitions. Distributions satisfying the conditions are called the **core**.

The core can be displayed graphically using barycentric coordinates. When we talk about a distribution of the win n , we mean a triple (a, b, c) of numbers with $a + b + c = n$. As pairs of numbers are usually displayed in the plane, triples of numbers can be represented in three dimensions. The triples of numbers that sum up to n form a plane, and if we add the requirement that a, b , and c should be positive, we have a triangle. We take it out of 3-dimensional space and put it into a plane, as in Figure 35.2 for $n = 10$. Triples (a, b, c) with nonnegative coordinates and $a + b + c = 10$ correspond to points in the triangle. The first coordinate expresses the distance of the point to the side on the bottom, the second coordinate is the distance to the upper right side, and the third coordinate is the distance to the upper left side.

The core for $n = 10$ and $m = 4$ is the shaded area in Figure 35.2. The requirements $a \geq 2, b \geq \frac{10}{3}, c \geq \frac{2}{3}$ mean that the points must be beyond the green lines, and the requirements $a + b \geq \frac{20}{3}, a + c \geq 3, b + c \geq 5$ are indicated by the blue straight line borders.

35.4.2 The Shapley Value

Sometimes the core contains no triple, often it contains many triples. Another concept always delivers a unique distribution, the so-called **Shapley value**, which we now illustrate in our example.




Figure 35.2. All distributions of the grand coalition, the core, and the Shapley value

Assume that the grand coalition is built gradually. Let’s say Ann starts, then Beth joins to form a coalition with Ann, and then Cindy joins to form the grand coalition. Assume that when a player joins the current coalition, every player gets the surplus the larger coalition now expects. That means that Ann starts with an expectation of 2. When Beth joins, the coalition’s expectation raises to $\frac{20}{3}$, and therefore Beth brings $\frac{14}{3}$ to the table. Finally, when Cindy joins, the expectation raises to 10, so Cindy claims the difference of $\frac{10}{3}$ for herself.

This procedure is not fair. When Beth would start and Ann would join, Ann could claim a surplus of $\frac{10}{3}$ instead of the 2 she gets when starting the coalition. For this reason, the procedure is performed for every ordering of the players, and then the average over all orderings is taken. For three players there are six orderings, and the claims for the players are

Ann starts with 2	Beth adds $\frac{14}{3}$	Cindy adds missing $\frac{10}{3}$
Ann starts with 2	Cindy adds 1	Beth adds missing 7
Beth starts with $\frac{10}{3}$	Ann adds $\frac{10}{3}$	Cindy adds missing $\frac{10}{3}$
Beth starts with $\frac{10}{3}$	Cindy adds $\frac{5}{3}$	Ann adds missing 5
Cindy starts with $\frac{2}{3}$	Ann adds $\frac{7}{3}$	Beth adds missing 7
Cindy starts with $\frac{2}{3}$	Beth adds $\frac{13}{3}$	Ann adds missing 5

The averages of the values, which are the Shapley values, are 3.28 for Ann, 4.94 for Beth, and 1.78 for Cindy. This corresponds to the point shown inside the core in Figure 35.2. In our example, it lies near the center of the core, but this does not always happen.

Exercises

- 1. Analyze the case where there is a coalition of Ann and Cindy, splitting win and loss 50:50, for $n = 10$, and $m = 4$. Draw the game tree and perform a backward induction analysis.
- 2. Analyze the case where there is a coalition of Ann and Beth, splitting win and loss 50:50, for $n = 10$, and $m = 4$. Draw the game tree and perform a backward induction analysis.
- 3. Analyze the case where there is a coalition of Beth and Cindy, splitting win and loss 50:50, for $n = 12$, and $m = 4$. Draw the game tree and perform a backward induction analysis.

Project 59

Analyze the negotiation process with coalitions where the split can be negotiated for the parameters $n = 10$, $m = 4$. Show that there is no stable coalition. Can you predict what might happen in the negotiations?

CHAPTER 36

Example: VNM POKER(4, 4, 3, 5)

Prerequisites: Chapter 25 and all theory chapters except maybe Chapter 32.

This chapter demonstrates the difficulties one may face when looking for mixed Nash equilibria in larger examples. Moreover we see that though one player’s Nash equilibrium mix may draw out many more strategies than the other player’s corresponding Nash equilibrium mix, playing the mix may still be worthwhile in two-player zero-sum games.

Class Activity In the applet [VNMPoker4](#), select values m and n , one of the computer opponents, and play 30 rounds.

This example is more complex than the case $S = 2$ discussed in Chapter 31. An analysis of the family of games gets too complex for general parameters r , m , and n , so we choose $r = 4, m = 2$, and $n = 3$.

36.1 Mixed Nash Equilibria

Pure strategies for Ann are four-letter words of Cs and Rs, and Beth’s pure strategies are four-letter words made up of Fs and Cs. The normal form is calculated using the Excel sheet [VNMPoker4.xlsx](#). We showed in Chapter 25 that some pure strategies are weakly dominated. After eliminating them we got the matrix in Table 36.1

Applying Brown’s fictitious play, and running, say, 1000 rounds in the Excel file [Brown10.xlsm](#), we can be fairly sure that Ann should play strategy CCCR 3/4 of the time and RCCR 1/4 of the time. Beth’s result is less clear. According to the results I obtained, Beth should mix about 5% of FFFC, about 37% of FFCC,

	FFFC	FFCC	FCCC	CCCC
CCCR	0	$\frac{2}{15} \approx 0.133$	$\frac{4}{15} \approx 0.267$	$\frac{2}{5} = 0.4$
CCRR	$\frac{1}{60} \approx 0.017$	0	$\frac{4}{15} \approx 0.267$	$\frac{8}{15} \approx 0.533$
CRCR	$\frac{5}{12} \approx 0.417$	$\frac{1}{60} \approx 0.017$	0	$\frac{4}{15} \approx 0.267$
CRRR	$\frac{13}{30} \approx 0.433$	$\frac{-7}{60} \approx -0.117$	0	$\frac{2}{5} = 0.4$
RCCR	$\frac{49}{60} \approx 0.817$	$\frac{5}{12} \approx 0.417$	$\frac{1}{60} \approx 0.017$	0
RCRR	$\frac{5}{6} \approx 0.833$	$\frac{17}{60} \approx 0.283$	$\frac{1}{60} \approx 0.017$	$\frac{2}{15} \approx 0.133$
RRCR	$\frac{37}{30} \approx 1.233$	$\frac{3}{10} = 0.3$	$\frac{-1}{4} = -0.25$	$\frac{-2}{15} \approx -0.133$
RRRR	$\frac{5}{4} = 1.25$	$\frac{1}{6} \approx 0.167$	$\frac{-1}{4} = -0.25$	0

Table 36.1. VNM POKER(4, 4, 3, 5), with some weakly dominated strategies eliminated

	FFFC	FFCC	FCCC	CCCC
CCCR	0	$\frac{2}{15} \approx 0.133$	$\frac{4}{15} \approx 0.267$	$\frac{2}{5} = 0.4$
RCCR	$\frac{49}{60} \approx 0.817$	$\frac{5}{12} \approx 0.417$	$\frac{1}{60} \approx 0.017$	0

Table 36.2. VNM POKER(4, 4, 3, 5), with more strategies eliminated

and about 58% of FCCC. Fortunately, all we need to know is that Ann mixes only CCCR and RCCR to find her optimal solution. We get the matrix in Table 36.2:

We now use the graphical method for solving $2 \times n$ zero-sum games. We draw four straight lines: The FFFC line from $(0, 0)$ to $(1, \frac{49}{60})$, the FFCC line from $(0, \frac{2}{15})$ to $(1, \frac{5}{12})$, and so on. The height of the FFFC line at p indicates Ann’s payoff if Ann plays RCCR with probability p and CCCR with probability $1 - p$, and if Beth plays FFFC, and similarly for the other three lines. See Figure 36.1. Since Beth wants to maximize her payoff, and therefore wants to minimize Ann’s payoff (the game is zero-sum), for each such p Beth would respond with the pure strategy that has the lowest height. According to the graph, for $0 \leq p \leq \frac{1}{4}$ (a lot of CCCR), Beth would play FFFC; for $\frac{1}{4} \leq p \leq 0.89$, Beth would play FCCC; and otherwise Beth would play CCCC. The piecewise-linear curve on the top of the gray area indicates the payoff Ann can achieve when playing a mix of p of RCCR and $1 - p$ of CCCR. Since Ann wants to maximize her payoff, she chooses that p where the curve has the largest height, so she chooses $p = \frac{1}{4}$. The expected payoff for Ann is $\frac{49}{240} \approx 0.2$.

It is a coincidence that the lines FFFC, FFCC, and FCCC intersect at a point. That means that Beth can mix the three strategies. To find the percentages q_1, q_2 , and $q_3 = 1 - q_1 - q_2$, we again use the Indifference Theorem and a little algebra.




Figure 36.1. Graphical solution of the 2×4 case

Ann’s best responses to Beth’s mix of q_1 of FFFC, q_2 of FFCC, and $1 - q_1 - q_2$ of FCCC are CCCR and RCCR. We equate Ann’s payoff for them:

$$0 \cdot q_1 + \frac{2}{15} \cdot q_2 + \frac{4}{15} \cdot (1 - q_1 - q_2) = \frac{49}{60} \cdot q_1 + \frac{5}{12} \cdot q_2 + \frac{1}{60} \cdot (1 - q_1 - q_2).$$

The equation does not have a unique solution. We get a variety of solutions, and so a variety of Nash equilibria mixes for Beth, all counterparts to the Ann’s mix.

The equation becomes

$$8 \cdot q_2 - 16 \cdot q_2 - 25 \cdot q_2 + q_2 = -16 + 1 + 16 \cdot q_1 + 49 \cdot q_1 - q_1$$

or

$$q_2 = \frac{15}{32} - 2 \cdot q_1.$$

The relationship between q_1 and q_3 is therefore

$$q_3 = 1 - q_1 - q_2 = 1 - q_1 - \frac{15}{32} + 2 \cdot q_1 = \frac{17}{32} + q_1.$$

Since probabilities cannot be negative, q_2 must be between 0 and $\frac{15}{32} \approx 47\%$, q_1 must be between 0 and $\frac{15}{64} \approx 23\%$, and q_3 must be between $\frac{17}{32} \approx 53\%$ and $\frac{49}{64} \approx 77\%$.

The relationship between q_2 and q_1 is linear, and so is that between q_3 and q_1 . In Figure 36.2, q_2 , q_3 , and q_1 are displayed as functions of q_1 . Since the relationships are linear, their graphs are straight lines. Every vertical line indicates a Nash equilibrium mix. The one shown has $q_1 = 6\%$, $q_2 = 34\%$, and $q_3 = 59\%$.




Figure 36.2. Relationship between probabilities q_1 , q_2 , q_3 for FFFC, FFCC, and FCCC in Beth's Nash equilibria

36.2 Performance of Pure Strategies against the Mixed Nash Equilibria

Playing a Nash equilibrium strategy against any of the other player's strategies in a zero-sum game guarantees a certain expected payoff. It is the value of the game. Unfortunately, the Nash equilibrium (mixed) strategy will not have a higher expectation against many of the other player's other strategies. All pure strategies that occur with nonzero probability in the other player's Nash equilibrium mixed strategy yield the same expected payoff. So, against sophisticated play, some less sophisticated play is sometimes not punished. In VNM POKER(2, r , m , n) for instance, all a player needs to know when playing against the Nash equilibrium is to avoid folding or checking with the higher value card. What about VNM POKER(4, 4, 3, 5)?

Let us see how Beth's pure strategies perform against Ann's optimal mix of $\frac{3}{4}$ CCCR and $\frac{1}{4}$ RCCR. We need the values in the full normal form in [VNMPoker4.xlsx](#) for the calculations. See the tab called "some mixes".

- FFFC, FFCC, FCFC, and FCCC are optimal against Ann's Nash mix. If Beth mixes them in any way, she will expect to lose only $\frac{49}{240} \approx 0.2$
- CFFC, CFCC, CCFC, and CCCC have an expectation of about -0.3 for Beth when played against Ann's Nash mix.
- FFFF, FF CF, FCFF, and FCCF have an expectation of about -0.49 against Ann's Nash mix.
- CFFF, CF CF, CCFF, and CCCF are Beth's weakest pure strategies, since they yield an expectation of about -0.58 for her.

That strategy FCFC that we eliminated as weakly dominated is still optimal here. Against CCCR and RCCR it behaves exactly as FFCC, so instead of playing FFCC with probability q_2 , we could play FFCC and FCFC with probabilities summing to q_2 . This further increases the number of Nash equilibria.

Let’s see how Ann’s pure strategies perform against Beth’s Nash mixes of q_1 of FFFC, $\frac{15}{32} - 2 \cdot q_1$ of FFCC, and $\frac{17}{32} + q_1$ of FCCC, for q_1 varying from 0 to $15/64 \approx 0.234375$. We don’t include FCFC here, since this would further complicate matters. The result is shown in Figure 36.3, based on the tab “some mixes” in the Excel sheet [VNMPoker4.xlsx](#). The graph displays Ann’s payoff versus q_1 .




Figure 36.3. Ann’s payoffs using pure strategies against Beth’s Nash mixes, varying q_1 from 0 (left) to 0.234375 (right)

For most cases, CCCR and RCCR are best responses for Ann. But to the right of the chart, the curves actually cross, and CCRR and RCRR are better responses to Beth’s mix. This holds for $q_1 > \frac{15}{68} \approx 0.2206$. That implies that the cases with q_1 between 0.2206 and 0.234375 do not form Nash equilibria. Maybe concentrating only on mixes of CCCR and RCCR for Ann was a mistake, so let us see whether mixes of CCRR and RCRR could form a Nash equilibrium for Ann. Then the matrix of Ann’s payoffs is

	FFFC	FFCC	FFCC	CCCC
CCRR	$\frac{1}{60} \approx 0.017$	0	$\frac{4}{15} \approx 0.267$	$\frac{8}{15} \approx 0.533$
RCRR	$\frac{5}{6} \approx 0.833$	$\frac{17}{60} \approx 0.283$	$\frac{1}{60} \approx 0.017$	$\frac{2}{15} \approx 0.133$

Table 36.3. VNM POKER(4, 4, 3, 5), with other strategies eliminated

We obtain the graph in Figure 36.4

The highest point of the gray area is exactly in the middle of the graph, so Ann mixes 50% of CCRR and 50% of RCRR. We can conclude that Beth mixes $\frac{15}{32}$ of FFCC and $\frac{17}{32}$ of FCCC.

Is this another Nash equilibrium? No, since Ann’s best responses to Beth’s mixing $\frac{15}{32}$ of FFCC and $\frac{17}{32}$ of FCCC are again only CCCR and RCCR. So assuming a Nash equilibrium in which Ann mixes CCRR and RCRR results in a contradiction and is therefore not possible.

Summarizing, we have found infinitely many Nash equilibria, Ann mixing 75% of CCCR and 25% of RCCR, and Beth mixing q_1 of FFFC, $\frac{15}{32} - 2 \cdot q_1$ of FFCC or FCFC, and $\frac{17}{32} + q_1$ of FCCC, for $0 \leq q_1 = \frac{15}{68} \approx 0.2206$. For large q_1 the percentage for FCFC is probably limited further. Ann’s payoff in the Nash equilibria is $\frac{49}{240} \approx 0.2$. Since we made assumptions on what pairs of strategies are used in Ann’s Nash mix, we cannot exclude the possibility of further Nash equilibria.




Figure 36.4. Analysis of the 2×4 case

We can draw another conclusion from the graph: since most of Ann’s pure strategy payoffs against the mixes are smaller for small q_1 than for larger q_1 , a clever Beth might choose $q_1 = 0$ to profit more from Ann making mistakes and not mixing CCCR and RCCR. Accordingly Beth might mix $\frac{15}{32}$ of FFCF and $\frac{17}{32}$ of FCCC. This Nash equilibrium might seem more reasonable than the other Nash equilibria.

It follows that if Ann plays the Nash mix mentioned, she will not gain any advantage if Beth plays a mix of FFCF, FFCC, FCFC, and FCCC. This translates into a behavioral strategy of F??C. Beth should always fold with a lowest value, and always call with a highest value, but otherwise she can do whatever she wants. A mixed strategy that sometimes calls with the lowest value and sometimes folds with the highest value will perform sub-optimally against Ann’s optimal mix. On the other hand, Beth’s playing any Nash equilibrium mix (with q_1 strictly smaller than $\frac{15}{68}$) will not gain her an advantage against Ann’s using a mix between CCCR and RCCR. Such a mix translates into the behavioral strategy ?CCR: always checking with a value of 2 or 3, always raising with a highest value of 4, but doing anything with a lowest value of 1. However, every mixed strategy that sometimes raises with a value of 2 or 3, or sometimes checks with a value of 4, will perform sub-optimally against Beth’s optimal mix of FFCF and FCCC. Decide for yourself whether you would have chosen any of the sub-optimal mixes. In the applet [VNMPoker4](#), for $m = 3$ and $n = 5$, all opponents play sub-optimally either as Ann or Beth except Ali Baba, Jim Knopf, U Hu, and I Can. Jim Knopf plays the Nash equilibria mixes discussed. U Hu, and I Can play Nash equilibria mixes for $m = 1, n = 2$ and $m = 1, n = 3$. In the applet [VNMPoker4CC](#), you can check the long-term performance of the strategies and some others provided by students. .

CHAPTER 37

Example: KUHN POKER(3, 4, 2, 3)

Prerequisites: All theory chapters, in particular chapters 24, 27, 32, and also Chapter 25.

Student Activity Play ten rounds of KUHNPOKER(3, 4, 2, 3) in the applet [KuhnPoker3](#) to refresh your familiarity with the game.

In Chapter 25 we gave a description of the games and provided extensive forms. We found that some pure and reduced pure strategies are weakly dominated and described the normal form for a few cases. Although some small cases of the games had pure Nash equilibria, we saw this is not true for KUHN POKER with $S = 3$. In this chapter we find a Nash equilibrium in mixed strategies for KUHNPOKER(3, 4, 2, 3). In our discussion we see why behavioral strategies in the game seem to be better than mixed strategies, and we also discuss their relation.

We saw in Chapter 25 that Ann’s first three information sets are her decision to raise or check with a jack, queen, or king, and her other three information sets occur after she has checked and Beth has raised, holding the same three cards. Beth’s first three information sets are her decision to call or fold when Ann has raised and Beth holds a jack, queen, or king, and Beth’s other three information sets occur when Ann has checked and Beth has to decide whether to check or raise with the same three cards. Some of Ann’s pure strategies can be reduced, since if she raises with some card, she will never have to think about whether she would fold or call it. After eliminating weakly dominated strategies, we arrive at the reduced normal form in Table 37.1.

	FFC	FFC	FFC	FFC	FCC	FCC	FCC	FCC	CCC	CCC	CCC	CCC
	CCR	CRR	RCR	RRR	CCR	CRR	RCR	RRR	CCR	CRR	RCR	RRR
CCCFFC	0.00	−0.06	−0.55	−0.61	0.00	−0.06	−0.55	−0.61	0.00	−0.06	−0.55	−0.61
CCC FCC	−0.12	0.00	−0.06	0.06	−0.12	0.00	−0.06	0.06	−0.12	0.00	−0.06	0.06
CCRFF●	0.00	−0.18	−0.67	−0.85	0.12	−0.06	−0.55	−0.73	0.24	0.06	−0.42	−0.61
CCRFC●	−0.12	−0.12	−0.18	−0.18	0.00	0.00	−0.06	−0.06	0.12	0.12	0.06	0.06
CRCF●C	0.06	0.18	0.00	0.12	−0.12	0.00	−0.18	−0.06	0.00	0.12	−0.06	0.06
CRRF●●	0.06	0.06	−0.12	−0.12	0.00	0.00	−0.18	−0.18	0.24	0.24	0.06	0.06
RCC●FC	0.55	0.48	0.18	0.12	−0.06	−0.12	−0.42	−0.48	−0.24	−0.30	−0.61	−0.67
RCC●CC	0.42	0.55	0.67	0.79	−0.18	−0.06	0.06	0.18	−0.36	−0.24	−0.12	0.00
RCR●F●	0.55	0.36	0.06	−0.12	0.06	−0.12	−0.42	−0.61	0.00	−0.18	−0.48	−0.67
RCR●C●	0.42	0.42	0.55	0.55	−0.06	−0.06	0.06	0.06	−0.12	−0.12	0.00	0.00
RRC●●C	0.61	0.73	0.73	0.85	−0.18	−0.06	−0.06	0.06	−0.24	−0.12	−0.12	0.00
RRR●●●	0.61	0.61	0.61	0.61	−0.06	−0.06	−0.06	−0.06	0.00	0.00	0.00	0.00

Table 37.1. KUHN POKER(3, 4, 2, 3), with some weakly dominated strategies eliminated

A slightly larger 18×16 matrix appears in the sheet “Start” in the file [KUH3.xlsm](#). In it you can automatically look for domination, pure Nash equilibria, and perform Brown’s fictitious play.

According to Nash’s Theorem, the game must have at least one Nash equilibrium in mixed strategies. Since it is a zero-sum game, it can be found by Brown’s fictitious play. In the Excel file, this can be found on the sheet called “Brown”. We get a pair of $\frac{2}{3}$ of $CCRFC\bullet$ and $\frac{1}{3}$ of $RCR\bullet C\bullet$ for Ann, and of $\frac{2}{3}$ of $FCCCCR$ and $\frac{1}{3}$ of $FCCRCR$ for Beth. The expected payoff for Ann is -0.02 , so KUHN POKER(3, 4, 2, 3) has a slight advantage for Beth.

How does one execute such a mixed strategy? In poker games, behavioral strategies are more natural for humans than mixed strategies. We need to be able to translate mixed strategies into behavioral strategies—this is done in Section 37.2. We also need to provide analysis going the other direction, since we should be able to describe expectations when the two players play behavioral strategies. This is easier, and we do it first, in Section 37.1.

37.1 From Behavioral Strategies to Mixed Strategies to Expectations

A behavioral strategy for Ann is a sequence of six numbers between 0 and 1, indicating the probabilities for raising in the first three information sets, or for calling in the other three information sets. Beth’s behavioral strategies are sequences of six numbers between 0 and 1 indicating her probabilities for calling and, respectively, raising.

Let’s translate the behavioral strategies into mixes of pure strategies. The method is straightforward. For a pure strategy, for instance $RCRFCC$ for Ann, we look at the six information sets and multiply the probabilities for the chosen options in the pure strategy as they appear in the behavioral strategy. In case of a $\bullet$ in a reduced strategy, we multiply by 1.

Let us illustrate it with Ann’s behavioral strategy (0.3, 0.8, 1, 0, 0.75, 1). When playing it, Ann always raises with a king, never calls with a jack, and always calls with a king. All pure strategies of the form $xyCzvw$, $xyzCvw$, and $xyzvwF$ are played with probability 0. For the other pure strategies, examples for the calculations are $p(CCRFF\bullet) = 0.7 \cdot 0.2 \cdot 1 \cdot 1 \cdot 0.25 \cdot 1 = 0.035$, $p(CRRF\bullet\bullet) = 0.7 \cdot 0.8 \cdot 1 \cdot 1 \cdot 1 \cdot 1 = 0.56$, and so on. The mixed strategy has nonzero probability only for the six reduced pure strategies $CCRFF\bullet$, $CCRFC\bullet$, $CRRF\bullet\bullet$, $RCR\bullet F\bullet$, $RCR\bullet C\bullet$, and $RRR\bullet\bullet\bullet$.

Beth’s behavioral strategy (0, 0.75, 1, 0.5, 0.8, 1) (never calling with a jack, always calling with a king, and always raising with a king) implies that only pure strategies of the form $FxCyZR$ occur with nonzero probability. Examples are $p(FFCRRR) = 1 \cdot 0.25 \cdot 1 \cdot 0.5 \cdot 0.8 \cdot 1 = 0.01$, $p(FCCCCR) = 1 \cdot 0.75 \cdot 1 \cdot 0.5 \cdot 0.2 \cdot 1 = 0.075$, and so on. The eight pure strategies $FFCCCR$, $FFCCRR$, $FFCRCR$, $FFCRRR$, $FCCCCR$, $FCCCCR$, $FCCRCR$, and $FCCRRR$ occur with nonzero probability in Beth’s mixed strategy.

On the sheet “Behavior” in the Excel file [KUH3.xlsm](#) the calculations are done automatically if you input at the top the behavioral strategies for the two players. The probabilities appear in the green and blue cells. The sheet does not use the full matrix because weakly dominated strategies have been removed. There are behavioral strategies that use some weakly dominated pure strategies with nonzero probabilities—in such cases you would need to use the sheet “BehFull”, which uses the full 27×64 matrix.

If we have found a mixed strategy for every behavioral strategy of our robots, it is fairly easy to calculate the expected payoffs—the sum of products of probabilities of an outcome and its payoff. Each pair of pure strategies, S_A for Ann and S_B for Beth, defines such an outcome, a cell in the payoff matrix corresponding to row S_A and column S_B , and the probability for that outcome, the probability for Ann choosing S_A and Beth choosing S_B is just the product of the probabilities of Ann choosing S_A in her mix, and Beth choosing S_B .

37.2 From Mixed Strategies to Behavioral Strategies

Which behavioral strategies correspond to the mixed strategies? To answer this more difficult question, we employ the methods in Chapter 32 of transforming mixed strategies into behavioral strategies.

The question is easier for Beth to answer, since her information sets don't have a history for Beth. In each of Beth's information sets Beth is about to move for the first time, and her probability is the mixed strategies probability of Beth's making that move. Beth would fold with a jack and call with queen or king when Ann has raised, and when Ann has checked, Beth would raise in $1/3$ of the cases with a jack, never with a queen, and always with a king. Therefore Beth's behavioral strategy would be $(0, 1, 1, \frac{1}{3}, 0, 1)$.

The same can be done for Ann in her first three information sets. Ann would accordingly raise in $1/3$ of the cases with a jack, never with a queen, but always with a king. The other three information sets of Ann have an information set history—Ann knows that she checked the move before. Therefore Ann would consider only those pure strategies in the Nash mix that do that. Then Ann's percentages for calling or folding are the percentages in the pure strategies that call or fold considered in the information set. For instance, in Ann's fourth information set, if Ann has a jack and has to call or fold, she would consider only $CCRFC\bullet$, since $RCR\bullet C\bullet$ would never result in this fourth information set. Therefore Ann folds. In the fifth information set, where Ann holds a queen, we consider the pure strategies $CCRFC\bullet$ and $RCR\bullet C\bullet$; since both always call with a queen, Ann will always call in the fifth information set. Finally, if Ann has a king and is about to call or fold, neither of the two pure strategies agrees with the information set history (of checking with a king). Therefore we consider no pure strategy, and we can choose any probability, since we will never face this information set. Therefore Ann's behavioral strategy is any of $\frac{1}{3}, 0, 1, 0, 1, x$ with x any number $0 \leq x \leq 1$. It is an interesting coincidence that Beth's recommendation is also one for Ann. For other parameters this is not true.

One word about bluffing, for some the most fun part of poker. In Kuhn poker, bluffing consists of raising with a jack, and both Ann and Beth can and should bluff occasionally. One should raise only with low (occasionally) and high (always) cards, but never for middle value cards. This is a pattern that von Neumann and Morgenstern observed in their analysis of variants of VNM POKER.

Usually, if we have a Nash equilibrium of Mix_A versus Mix_B , then Mix_B is not the only mixed strategy Beth has that behaves optimally against Mix_A . According to the Indifference Theorem, each of Beth's pure strategies that occurs with positive probability in Mix_B is a best response for Beth against Mix_A . Also, every mix of these strategies works. We can say more: even strategies that have been eliminated as weakly dominated may behave optimally against Mix_A . In our example, the only best responses to Ann's Nash equilibrium mix of $2/3$ of $CCRFC\bullet$ and $1/3$ of $RCR\bullet C\bullet$ are $FCCCCR$ and $FCCRCR$ of Beth's Nash mix, and also $FCCCR$ and $FCCRRR$. The only best responses to Beth's Nash mix are $FCCCCR$ and $FCCRCR$ of Ann's corresponding Nash mix.

Looking at these mixed strategies, and the behavioral strategies that can occasion them, we see that the only optimal behavioral strategies against the Nash mix have the form $(x, 0, 1, y, 1, z)$ with $y = x$ for Ann; and $(0, 1, 1, u, v, 1)$ for Beth.

Exercises

1. Find the mixed strategy corresponding to the behavioral strategy $\frac{1}{4}, 0, 1, 0, \frac{1}{2}, 1$ —raising with jack, queen, king in $\frac{1}{4}, 0$, or 1 of the cases, and calling with a jack, queen, or king in $0, \frac{1}{2}$, or 1 of the cases. for Ann.
2. What is Ann's behavioral strategy for the mixed strategy consisting of $\frac{1}{4}$ of $CCRFC\bullet$, $\frac{1}{4}$ of $RCR\bullet C\bullet$, $\frac{1}{4}$ of $CRRC\bullet\bullet$, and $\frac{1}{4}$ of $CCRFF\bullet$?
3. If Ann plays $RCR\bullet F\bullet$ and Beth plays $FCFRCR$, what is the payoff for Ann in the nine cases of card distributions?

4. If Ann plays $RCR \bullet F \bullet$ and Beth plays a 50% – 50% mix of $FCFR CR$ and $FCCCR R$, what is the expected payoff for Ann in the nine cases of card distributions?

CHAPTER 38

Example: End-of-Semester Poker Tournament

Prerequisites: Chapter 25 and all theory chapters, in particular Chapters 24, 27, and 32.

Each semester my game theory class has a poker tournament, playing either a version of KUHN POKER or a version of VNM POKER. The assignment looks like this.

Robot-KUHN POKER(3, 4, 2, 3) Tournament Every student will submit one poker-playing robot and will provide

- A nice sounding name for the robot,
- The three probabilities that dictate when the robot computer player, having the first move, should raise when facing the lowest value card jack J, the middle-valued card queen Q, and the highest value card king K.
- The three probabilities that dictate when the robot should call when facing a jack, a queen, or a king, having checked in the first move and having seen the other player raise.
- The three probabilities that dictate when the robot should call when facing a jack, a queen, or a king, having the second move and having seen the first player raise.
- The three probabilities that dictate when the robot should raise when facing a jack, a queen, or a king, having the second move and having seen the first player check.

Class Activity Create your own robot. What twelve numbers would you choose, and why? Discuss your choices with classmates.

In Spring 2011, the students and teacher submitted the robots in Table 38.1. Since the string of twelve numbers uniquely determines the robot’s behavior, we call it the DNA of the robot. To get sixteen players for a knockout tournament, three more robots were added: two identical Random robots, which always decide randomly with 50% probability, and one called SophRandom, which never folds or checks when facing a king but otherwise decides randomly with 50% probability. To make the outcome in our robot tournament depend more on the quality of the robot and less on luck, in each pairing the contestants played 200 rounds. The robot that increased its amount of money was declared the winner. You can find the robots and the playground in the applet [KuhnPoker3CC](#). Who do you think won the tournament. Who should have won?

Name	When having the first move (playing as Ann)						When having the second move (playing as Beth)					
	probability for raising			probability for calling			probability for calling			probability for raising		
	J	Q	K	J	Q	K	J	Q	K	J	Q	K
Bartleby	0.2	0.5	1	0.2	0.6	1	0	0.4	1	0.1	0.5	1
Hierophant	0	0.72	1	0	0.23	1	0	0.5	1	0	0.72	1
Voltron	0.3	0.7	1	0.1	0.6	1	0.2	0.4	1	0.4	0.8	1
Bill Nye	0.5	0	1	0.25	0.75	1	0.5	0.75	1	0	0.5	1
Yogi Bear	0.75	0.99	1	0	1	1	0	0.67	1	0.1	1	1
MP Hamster	0.75	1	1	0.5	1	1	0.45	1	1	0.3	0.8	1
Robo-Boogie	0.25	0.75	1	1	0.5	0.25	0.75	0.5	0.25	1	0.75	0.5
Amadeus	0.33	0	1	0	1	1	0	1	1	0.33	0	1
WN Turtle	0	0.75	1	0	0.75	1	0	0.75	1	0	0.75	1
Rupert	0.3	0.8	1	0	0.75	1	0	0.75	1	0.5	0.8	1
Venus	0.75	0.8	1	0.5	0.6	0.7	0.5	0.6	0.7	0.5	0.6	0.7
Jocker	0	0	0	0	1	1	0	1	1	0	0	1
Max	0	0.7	1	0	0.7	1	0	0.8	1	0	0.6	1

Table 38.1. The robots and their DNA

38.1 Expectations

The first six numbers (the second six numbers) of the DNA form behavioral strategies for Ann (for Beth). For every information set, probabilities for the options are given. Figure 38.1 displays the expected number of points for each pairing (for playing 200 rounds, with alternating first moves). We explain in Chapter 37 how to compute the numbers. Cells with positive value are in bold face. The rows have been sorted according to the numbers in the next to last column, which is the average of all values in the row, i.e., the robot’s average payoff when it plays against the other robots. The last column counts the number of wins in each row, i.e., the number of other robots that this robot would beat. There seems to be a mild correlation between a robot’s average payoff and its number of wins.

Expectations were not always met by reality. For instance, Hierophant played against Bill Nye in the first round, and although the expected outcome of the 200 round game was 16.22, the actual outcome was −10, meaning that Bill Nye won. So we see a deviation of −26,22 here. The deviations of the other plays in the first and second round of the knockout tournament were −13.6, −46.2, −36.9, 52.8, −84.7, −71.5, 2.4, −32.1, and −26.5. Thus, although larger deviations occurred, the typical deviation was between −50 and 50. This suggests that if the expectations are higher than 150 or so (as in the 167 payoff from the 200 rounds Rupert played versus the Random robot), the experiment went as expected.

Amadeus is the robot corresponding to the Nash equilibrium mix. We see that Amadeus has a nonnegative expectation against every other robot in the symmetric game consisting of an even number of rounds alternating as Ann and Beth. In Chapter 37 we saw that the only optimal behavioral strategies against the Nash Mix (the Amadeus robot), the only ones achieving an expected payoff of 0 against Amadeus, have the DNA of $x, 0, 1, y, 1, z | 0, 1, 1, u, v, 1$ with $y = x$. In the list, none of the robots except Amadeus has this form. Therefore, although in the long run many robots could possibly tie against the Nash mix, none showed up in the experiment .

	MP Hamster	Yogi Bear	Rupert the robot	Amadeus	WN Turtle	Max	Voltron	SophRandom	Jocker	The Hierophant	Bill Nye	Bartleby	Venus	Random1	Robo-Boogie	Average	Number of Wins
MP Hamster	0.0	20.7	6.1	-15.1	3.9	0.6	52.4	64.4	-10.9	36.2	38.6	55.6	121.4	236.8	260.5	58	12
Yogi Bear	-20.7	0.0	5.5	-13.3	20.1	16.1	52.3	54.2	-10.3	49.8	32.2	61.8	97.8	215.2	265.5	55	11
Rupert the robot	-6.1	-5.5	0.0	-11.3	25.0	24.9	29.7	38.5	-3.6	45.3	36.8	42.3	63.7	167.0	176.4	42	10
Amadeus	15.1	13.3	11.3	0.0	11.1	10.3	20.6	25.3	16.2	17.4	15.7	18.9	52.1	103.0	104.0	29	14
WN Turtle	-3.9	-20.1	-25.0	-11.1	0.0	0.7	-7.6	8.0	7.6	1.6	25.4	3.7	33.9	105.7	119.9	16	9
Max	-0.6	-16.1	-24.9	-10.3	-0.7	0.0	-8.9	6.0	11.9	-0.5	24.5	1.8	33.8	98.8	109.8	15	7
Voltron	-52.4	-52.3	-29.7	-20.6	7.6	8.9	0.0	6.1	-2.4	25.8	16.7	21.9	16.1	126.1	138.9	14	9
SophRandom	-64.4	-54.2	-38.5	-25.3	-8.0	-6.0	-6.1	0.0	0.0	15.0	6.8	17.5	14.8	124.2	121.2	6	6
Jocker	10.9	10.3	3.6	-16.2	-7.6	-11.9	2.4	0.0	0.0	-12.0	-9.1	-9.5	13.8	33.3	36.4	3	7
The Hierophant	-36.2	-49.8	-45.3	-17.4	-1.6	0.5	-25.8	-15.0	12.0	0.0	16.2	-2.7	-1.6	75.8	87.3	0	5
Bill Nye	-38.6	-32.2	-36.8	-15.7	-25.4	-24.5	-16.7	-6.8	9.1	-16.2	0.0	-1.3	13.9	85.6	78.8	-2	4
Bartleby	-55.6	-61.8	-42.3	-18.9	-3.7	-1.8	-21.9	-17.5	9.5	2.7	1.3	0.0	-7.7	73.6	92.6	-3	5
Venus	-121.4	-97.8	-63.7	-52.1	-33.9	-33.8	-16.1	-14.8	-13.8	1.6	-13.9	7.7	0.0	145.4	160.0	-10	4
Random1	-236.8	-215.2	-167.0	-103.0	-105.7	-98.8	-126.1	-124.2	-33.3	-75.8	-85.6	-73.6	-145.4	0.0	-33.3	-108	0
Robo-Boogie	-260.5	-265.5	-176.4	-104.0	-119.9	-109.8	-138.9	-121.2	-36.4	-87.3	-78.8	-92.6	-160.0	33.3	0.0	-115	1

Figure 38.1. Expected number of won/lost points, 200 rounds

38.2 Odds

But wait a minute! A poker tournament is not the same as an ordinary night of poker games! Since no money is transferred, a win of 80:−80 or a win of 5:−5 have the same payoff for the player—advancing to the next round. So it is not how much money is won that is important, but the odds—the probabilities p_A for a win after 200 rounds, p_0 for a draw, and p_B for a loss. If a win counts 1, a draw 0, and a loss −1, then Ann’s expected payoff is $p_A \cdot 1 + p_0 \cdot 0 + p_B \cdot (-1) = p_A - p_B$. Of course Beth’s expected payoff is $p_B - p_A$.

It seems likely that the Nash equilibria for the modified game of playing 200 rounds and having payoffs of 1, 0, or −1 at the end differ from the Nash equilibria of the winning money game discussed in the previous section. We will not try to see how but only show how to calculate p_A , p_0 , p_B , and the expected payoffs for pairs of strategies used by the students in the example above.

The expected payoff for 200 rounds of Amadeus versus Max is 10.3, but what are the odds for Amadeus winning? Typical deviations in 200 rounds range between −50 and 50. We will calculate the probabilities of an Amadeus win, a draw, and a Max win, in the 200 rounds of play. This will demonstrate how to make these calculations for every pairing.

What we need first are the probabilities that a pure Ann strategy like $CCRFC\bullet$ wins 3 units, 2 units, nothing, −2 units, and −3 units, against a pure Beth strategy like $FFCCCR$. We separately consider the nine card distributions to Ann and Beth and we check what the outcome will be if both players play according to their pure strategies. In our example, $CCRFC\bullet$ versus $FFCCCR$, Ann’s payoffs in the 0 cases are:

	J	Q	K
J	0	−2	−2
Q	2	0	−3
K	2	2	0

Therefore

- $CCRFC\bullet$ never wins 3 units.
- $CCRFC\bullet$ wins 2 units when Ann has a queen and Beth a jack, or Ann has a king and Beth a jack or a queen. Accordingly the probability for Ann’s winning 2 units is $3 \cdot p_{XY} = 3 \cdot 0.121 = 0.364$.

- *CCRFC*• draws in the cases where both have identical cards (which is not obvious, but depends on the strategies!). The probability for this is $3 \cdot p_{XX} = 3 \cdot 0.091 = 0.273$.
- *CCRFC*• loses 2 units for Ann if Ann has a jack and Beth a queen or king, which occurs with probability $2 \cdot p_{XY} = 2 \cdot 0.121 = 0.242$.
- *CCRFC*• loses 3 units for Ann if Ann has a queen and Beth a king, which occurs with probability $p_{XY} = 0.121$.

The calculations have to be done for all pairs of pure strategies. Then the probability for Ann to win 3 units if Ann uses a mixed strategy Mix_A and Beth a mixed strategy Mix_B is the obtained by looking at all pairs S_A and S_B of pure strategies, calculating the probability of this pair being chosen, which is the product of the occurrence probability of S_A in Mix_A and the occurrence probability of S_B in Mix_B , and multiplying this product by the probability for S_A winning 3 units against S_B . Then the products are added.

The Excel sheet [KUH3Odds.xlsx](#) automates the calculations. On the “Start” sheet, the parameters r, m , and n have to be entered into the black cells. Sheet “nm0-m-n” calculates the five probabilities (Ann winning $n, m, 0, -m, -n$) for all pairs of strategies. Then enter or copy and paste behavioral strategies into the black rectangles on sheet “Odds23”. In sheet “AB” the five probabilities $a_3, a_2, a_0, a_{-2}, a_{-3}$ for the left robot winning 3, 2, 0, -2, and -3 units are calculated if it plays Ann and the one on the right plays Beth. The values are displayed in the red cells. Similarly, in the sheet BA the probabilities $b_3, b_2, b_0, b_{-2}, b_{-3}$ are calculated if the right robot plays Ann and the left robot plays Beth.

In our example, Amadeus versus Max, if Amadeus plays the Ann role, Amadeus wins 3, 2, 0, -2, and -3 with probabilities $a_3 = 0.097, a_2 = 0.305, a_0 = 0.242, a_{-2} = 0.162$, and $a_{-3} = 0.194$. So Amadeus wins with probability $0.097 + 0.305 = 0.402$, draws with probability 0.242, and loses with probability $0.162 + 0.194 = 0.356$ when Amadeus moves first. If Max plays the Ann role, then Max wins 3, 2, 0, -2, and -3 with probabilities $b_3 = 0.130, b_2 = 0.230, b_0 = 0.242, b_{-2} = 0.287$, and $b_{-3} = 0.110$.




Figure 38.2. A probability digraph

What happens in two rounds, where Amadeus moves first in the first round (has the Ann role), and Beth moves first in the second round? This is a multi-step experiment that is best described by the probability digraph in Figure 38.2. The probabilities for the payoffs for Amadeus can be expressed in terms of a_i and b_j as shown in the digraph. The probabilities for Amadeus winning 6, 5, ..., -5, -6 units in the two rounds




Figure 38.3. Probabilities for winning or losing different amounts in two rounds for Amadeus versus Max

are 0.011, 0.062, 0.088, 0.050, 0.144, 0.040, 0.209, 0.095, 0.095, 0.078, 0.037, 0.066, and 0.025. They are displayed in the chart in Figure 38.3. The probability for Amadeus winning in two rounds is the sum of the probabilities of the states 1, . . . , 6, which is 39.4% in our example. It is the sum of the areas of the six rectangles on the left. The probability for a draw in two rounds is 20.9%, and the probability for Max winning is 39.7%, which is the sum of the areas of the six rectangles on the right side of the graph. Although Amadeus has a higher expectation in the two-round game, Max is a tiny little bit more likely to win!

But wasn't Amadeus supposed to beat or tie every other robot? Only in the expected value of money won. Although the chart has more area to the right, the area to the left extends more from the middle, so that the balancing point of the figure is a little left of middle.

We can extend the probability digraph to more than two rounds. The extensions are implemented on the sheet "Odds23" in the KUHN3Odds.xlsx file. When you input behavioral strategy values into the black rectangles in the blue area, the winning, drawing, and losing probabilities for round numbers between 1 and 250 are shown on the left border (cells B, C, and D). The chart in Figure 38.4 was created by the Excel sheet and shows the probabilities of Amadeus winning, drawing, or losing against Max, for different numbers of rounds. As the number of rounds increases, the probability for Amadeus winning exceeds the probability for his losing, and also increases. What do you think about this question: will the probability for Amadeus winning ever hit 70%, for some large number of rounds?




Figure 38.4. Probabilities for Amadeus winning, drawing, or losing against Max, for increasing number of rounds

38.2.1 Many Rounds

“The longer you play, the less luck is a factor. A single hand of poker is about 99 percent luck. After eight hours of playing, it is about 88 percent luck.” This is Chris Ferguson’s commentary on luck in poker [PP2005]. Although the numbers are surely not meant to be precise, isn’t it interesting how small the estimate of the skill factor in poker is, even after eight hours of play, by a professional, who is well-known for his analytical approach? What if we played longer than that? Could the influence of luck be reduced to almost zero?

Let me explain things with a simpler example, a so-called Bernoulli experiment. Assume a coin falls head with 51% probability and tail with 49%. Ann wins \$1 from Beth if heads shows, and otherwise Ann has to give \$1 to Beth. The expected payoff for Ann in one round is $0.51 \cdot 1 + 0.49 \cdot (-1) = 0.02$. Then in 100 rounds, Ann is expected to win \$2, in 400 rounds \$8, and in 1600 rounds \$32 dollars. In reality the payoffs will deviate. In 100 rounds, the deviation from the expected value typically lies between -5 and 5 (about 70% of the cases the deviation lies between these numbers). In 400 rounds, the deviation is larger, typically between -10 and 10 . The sizes of the deviations grow less than the expected values. In 1600 rounds, the typical deviations double again and lie between -20 and 20 , but since the expected value increased by a factor of 4 to 32, in more than 70% of the cases Ann will have won between \$12 and \$52. If we proceed to larger and larger numbers of rounds, we can increase the winning probability for Ann as close to 1 as we desire.

The same pattern occurs in our more complicated case. Provided a robot has a positive expected payoff against another player, we can force its winning probability to be close to 1 by increasing the number of rounds played. Not only does the curve above hit 70%, it will approach 100%, if we increase the number of rounds sufficiently.

38.3 The Favorite in Knockout Tournaments

Let’s return to our question at the end of Section 1: Who should have won? If we were to play a very large number of rounds, say a million, the results from Section 4 imply that the Nash mix robot, Amadeus, would win with very high probability against each of the other robots, since it has a positive expectation. Therefore, Amadeus is the favorite. However, for a smaller number of rounds such as 200, this doesn’t have to be the case. It depends on the population of the robots in the tournament and on chance.

Let us look at a smaller case—a knockout tournament of four players, Amadeus, Rupert, and two clones of Max, called Max1 and Max2. If there is a tie after 200 rounds, then 200 more rounds are played. Therefore odds of 56,3% : 1.3% : 42.3% for Ann winning, drawing, losing, may change into odds of 57% : 43% for winning and losing. Although this is not mathematically precise, to simplify matters we divide the probabilities proportional to the probabilities for winning and losing. The probabilities for the row player winning in the 200 rounds version are

	Amadeus	Max1	Max2	Rupert the robot
Amadeus	-	57%	57%	57.3%
Max1	43%	-	50%	33.4%
Max2	43%	50%	-	33.4%
Rupert the robot	42.7%	66.6%	66.6%	-

Table 38.2. Winning probabilities for the row robot against the column robot in 200 rounds

How likely is winning the tournament? There are three equally likely pairings: Amadeus-Max1 and Max2-Rupert, Amadeus-Max2 and Max1-Rupert, and Amadeus-Rupert and Max1-Max2.

- If the pairing is Amadeus-Max1 and Max2-Rupert, to win the tournament Amadeus must beat Max1 in the first round, and then beat whoever remains. The probability for beating Max1 in the first round is 57%.

Having achieved this, the probability that the opponent is Max2 and is defeated by Amadeus as well is $33.4\% \cdot 57\%$, and the probability that Rupert enters the final and is defeated by Amadeus is $66.6\% \cdot 57.3\%$. Therefore the probability for Amadeus to win the tournament is $57\% \cdot (33.4\% \cdot 57\% + 66.6\% \cdot 57.3\%) = 32.6\%$.

- In the same way, the probabilities for a victory of Max1, Max2, and Rupert in this pairing are 16.7%, 15.4%, and 35.3% respectively. Similarly we get 32.6%, 15.4%, 16.7%, and 35.4% as winning probabilities for Amadeus, Max1, Max2, and Rupert for the first round pairing Amadeus-Max2 and Max1-Rupert.
- If the first round pairing is Amadeus-Rupert and Max1-Max2, then the victory probabilities for Amadeus, Max1, Max2, and Rupert are 32.7%, 19.5%, 19.5%, and 28.4%.

For the total probabilities, we take averages. We get 32.6% for Amadeus, 17.2% for both Max1 and Max2, and 33% for Rupert. Rupert has the highest probability of winning the tournament! The explanation for this is that Rupert has better chances against the Maxes than Amadeus. Rupert's advantage over the two Maxes is larger than his disadvantage against Amadeus.

For knockout tournaments of 8 or 16 players, the same analysis is possible. However, an 8 player tournament has $7 \cdot 5 \cdot 3 \cdot 1 \cdot 3 = 315$ different pairing patterns for the first and second round, which makes analysis without the help of computers tedious.

38.4 Similarity of the DNA (optional)

The sequence of the twelve probabilities that create the behavioral strategies for a player is the DNA of a robot. Sequences could be close or far apart, depending on how many numbers are identical, or close, in the sequences. We could use the sum of the squares of the differences as a measure of closeness. For instance, the DNA of MP Hamster and Yogi Bear are close, the distance is

$$\begin{aligned}
 & (0.75 - 0.75)^2 + (1 - 0.99)^2 + (1 - 1)^2 + (0.5 - 0)^2 + (1 - 1)^2 + (1 - 1)^2 \\
 & + (0.45 - 0)^2 + (1 - 0.67)^2 + (1 - 1)^2 + (0.3 - 0.1)^2 + (0.8 - 1)^2 + (1 - 1)^2. \\
 & = 0 + 0.0001 + 0 + 0.25 + 0 + 0 + 0.2025 + 0.1089 + 0 + 0.04 + 0.04 + 0 = 0.642.
 \end{aligned}$$

This measure of similarity could be interpreted as a distance in 12-dimensional space. It is always between 0 and 12, with 0 meaning that both robots behave identically, whereas a distance of 12 would indicate that both play a pure strategy and both play the opposite in each information set. While one would as Ann always raise with a jack, the other would always check with a jack, and so on. For another example, the distance between the DNAs of MP Hamster and Jocker is 3.745.

It may be interesting to investigate whether closeness of the DNA implies similar performance against other robots. In the expected payoff table, there is similarity between Amadeus and Jocker, since both beat the first three robots in the list. The DNA of Amadeus is closest to the DNA of Jocker, with a distance of 1.22, whereas Robo-Boogie's DNA is farthest away with a value of 6.625. Closest to Robo-Boogie is the DNA of the random robot, having all twelve probabilities equal to 0.5. The robots are also close in performance.

38.5 How to Create your own Tournament

Maybe you want to create your own playground? Maybe not KUHN POKER(3, 4, 2, 3) or VNM POKER(4, 4, 3, 5), since these versions have already been analyzed? This is not difficult. First you decide which version to play. Possible are

- VNM POKER(4, r, m, n)
- VNM POKER(5, r, m, n)

- KUHN POKER(3, r , m , n),

for r , m , and n with $1 \leq r \leq 4$ and $1 \leq m < n \leq 7$. Then you collect the robot names and the behavioral strategies, the DNA from the (up to 16) participating persons. Each DNA is a string of 8, 10, or 12 numbers between 0 and 1. There are three templates for the playgrounds, the applets [VNMPoker4CC](#), [VNMPoker5CC](#), or [KuhnPoker3CC](#). Open the corresponding file using Notepad or any other text editor. Scroll down to the part that looks like

```
// -----
// CHANGE HERE FOR DIFFERENT PARAMETERS
// -----
// r:  (must be 1, 2, 3, or 4)
document.getElementById('rrr').value=4;
// m:
document.getElementById('lowbet').value=2;
// n:
document.getElementById('highbet').value=3;
```

Replace the values (here 4, 2, 3) for r , m , n by the values in your version. Below this is

```
// -----
// Change the Values Below
// -----
Name[0] = "Bartleby";
A1[0] = 0.2;
A2[0] = 0.5;
A3[0] = 1;
.....
// -----
// Change the Values Above
// -----
```

Replace the names and values by the names and DNA values of your robots. Save the file under a new name, with extension “html” in the “Applet” subfolder. It is now a web page that opens in your browser when you doubleclick it.

Exercises

1. If Ann plays $RCR \bullet F \bullet$ and Beth plays a 50%–50% mix of $FCFR$ and $FCRR$, what is the expected payoff for Ann in the nine card distributions? What is Ann’s probability of winning?
2. Select four robots (different from Amadeus, Max, and Rupert) from the list of robots, and calculate the probabilities for winning a 4-player knockout tournament (playing 200 rounds in each pairing) for each player.
3. Among the robots listed, find the robot whose DNA is closest to the Random robot who has a probability of $1/2$ in every position. Which of the robots is farthest away from the Random robot?

Project 60

Select five robots, different from Voltron, and calculate their similarity using the measure of sum of the squares of the differences of the twelve numbers in the DNA. Then compare these numbers with the behavior, the expected value, against Voltron. Do robots with similar DNA achieve similar results against Voltron? Try robots other than Voltron. Try to explain your findings.

Project 61

Calculate the average DNA of the thirteen robots listed by calculating in each information set the average of their probabilities. How does this average robot perform against the other thirteen robots? Would a robot that scores well against the average robot also score well in the tournament? Explain.

Project 62

Investigate the relationship between expected value and odds for two numbers of rounds as follows: Select ten pairings from the available robots (or invent your own).

1. For the ten pairings calculate the expected payoffs in a 200 round game, calculate the odds, and graph expected value versus odds for the player to the left winning. Is there a monotonic relationship, in the sense that larger expected value always implies higher winning probability?
2. Do the same for two-round games.
3. If there is any discrepancy between (a) and (b), try to explain it.

CHAPTER 39

Stockholm 1994

In Fall 1994, mathematician John F. Nash received a phone call that he had won a Nobel prize, together with economists John C. Harsanyi and Reinhard Selten. The next day the story filled the newspapers. The *New York Times* had a headline “Game theory captures a Nobel”. Mathematical journals reported the news weeks or months later.

Hold on a minute! There is no Nobel prize in mathematics! In the early twentieth century Alfred Nobel established prizes only in physics, chemistry, medicine, literature, and peace. Nash worked in none of these disciplines. His award was not one of the original Nobel prizes but was rather the Nobel Memorial Prize in Economic Sciences, which was initiated in 1968 and is sponsored by the Central National Bank of Sweden. Like the original Nobel prizes, it is awarded by the Royal Swedish Academy of Sciences. It is often through this prize that mathematicians hit the front pages of the newspapers and become famous.

Nowadays the academic areas closest to mathematics are probably physics, computer science, and economics. When Greek culture dominated the western world, philosophy was close to mathematics, but the two have separated since then. Astronomy and mathematics were closely related in ancient times, but the close relationship between physics and mathematics was established and strengthened in the 17th century through the development of calculus. Computer science started as a special branch of engineering and mathematics, and has been an independent area for not much longer than 40 years. The subfields of mathematical physics and theoretical computer science are still closely related to mathematics, and there are many exchanges between researchers working in both fields.

In contrast to the friendly kinships between the siblings physics and mathematics, and mathematics and computer science, the relationship between economics and mathematics has been more difficult. For two centuries there have been efforts to make economic theories more mathematical, and to describe economic problems with mathematical models. Some economists have been critical. They say that the models do not describe the economic problems well enough and that economic reality is too complex for the simple mathematical models to lead to meaningful results. Part of the criticism has to do with mathematical language which sometimes hides straightforward economic ideas under complicated mathematical terminology. Part of the criticism is the mantle of false authority some authors may try to claim by describing their ideas mathematically—can their ideas be wrong if their calculations are correct? Many recipients of the Nobel prize in economics have had Ph. D.s in mathematics, or written in a mathematical style, as Gerard Debreu, Kenneth Arrow, and Leonid Kantorovich.

Although John von Neumann and Oscar Morgenstern said in their famous book on game theory that they wanted to revolutionize economics, game theory was not used much in economics in the 50s and 60s. As Kuhn described it, “von Neumann’s theory was too mathematical for the economists.... As a consequence, the theory of games was developed almost exclusively by mathematicians in this period.” [K2004]. For the most part, classical economic models assume a huge number of rational agents who want to maximize their

payoffs given complete information. Exceptions are monopolistic models with just one seller. But there are models with more than one but still few agents, proceeding sometimes with incomplete (asymmetric) information, where players know different parts of the whole. Game theoretical models, did not become fashionable in economics until the 70s.

As mentioned, Reinhard Selten and John C. Harsanyi shared the 1994 Nobel prize with John Nash. Harsanyi was honored for his papers that showed that games of incomplete information can be transformed into games of complete but imperfect information. Selten introduced subgame perfect equilibria, showing that the extensive form has often more information than the normal form.

Once the spell was broken, many economists who have made major contributions to game theory have been awarded the Nobel prize. Many were mathematicians by training, and moved to departments of economics. I will list them.

- The 1996 Nobel prize was shared by James A. Mirrlees and William Vickrey for their work on incentives under asymmetric, incomplete information. Examples are auctions, or optimal taxation. Vickrey started auction theory in 1961 by introducing the procedure that is now called “Vickrey auction”. It is an auction with sealed bids, where the bidder with the highest bid gets the item, but pays only the second highest bid. Vickrey showed that this strange-looking auction has some desirable properties. For example, players will reveal their preferences by bidding what the item is worth to them.
- The 2001 Nobel prize was awarded to George A. Akerlof, A. Michael Spence, and Joseph E. Stiglitz for their work on games with incomplete (asymmetric) information. Akerlof had written a paper on how to formulate used car markets (where the buyer is unsure about the quality of the car) and similar situations as games. Spence formulated an education game, where high school graduates decide to attend college or not. College is easier for the capable, but getting a degree is possible for everybody. Spence claimed, provocatively, that it may not necessarily be the case that what students learned at college is important at their jobs—one needs only the assumption that those for whom college is easy (the capable) are also good at their job. So it is a game where capable persons have to convince others that they are capable, and they do that by going through college. This line of research started with concrete examples, with the idea of “beginning with very concrete problems and then generalizing them,” as Stiglitz put it in the Nobel prize interview [N2001]. Or, according to Akerlof in the same interview: “We argued from the bottom up. So we took a look at examples, such as insurance markets, and education markets, and credit markets, and markets for used cars, and then argued from the way we thought those specific examples worked to how general markets should look.”
- The 2005 Nobel prize was shared by Robert J. Aumann and Thomas C. Schelling. They are rather different: Aumann has a Ph. D. in pure mathematics, in knot theory, and writes in a mathematical style in his game theoretical papers. He also has a very “l’art pour l’art” view of academia, having said that “science is just about curiosity, about understanding the world, not necessarily of any practical use” in an interview in 2008 [N2008]. He was awarded the prize for his research on repeated games and cooperation. Schelling has a verbal style, and he has helped popularize many game theoretic ideas. His book *The Strategy of Conflict* [S1960] published in 1960, had an enormous influence.
- The 2007 Nobel prize went to Leonid Hurwicz, Eric S. Maskin, and Roger B. Myerson for work in mechanism design. The task of mechanism design is to create and establish games that yield desirable outcomes provided the players play optimally. Hurwicz laid the foundations around 1972.
- The 2012 Nobel prize went again, at least partially, for mechanism design. Alvin E. Roth shared it with Lloyd Shapley, a mathematician. Shapley and Roth had investigated matching systems, where agents (like students) are distributed to institutions (like colleges) according to a procedure that takes preferences of agents and institutions into account. Shapley had previously done work in game theory, mainly in cooperative game theory.

There are a number of eminent game theorists that did not get the Nobel prize, Harold William Kuhn, Martin Shubik, and David Gale, to name just a few.

Bibliography

- [A1989] J. D. Allen, A note on the computer solution of connect-four, in *Heuristic Programming in Artificial Intelligence*, D.N.L. Levy, D.F. Beal, eds., Ellis Horwood, 1989, 134–135.
- [A1988b] L. V. Allis, A knowledge-based approach to connect-four, the game is solved: white wins, Master's Thesis, Vrije Universiteit Amsterdam, 1988.
- [A1988] William Aspray, The emergence of Princeton as a world center for mathematical research, 1896–1939, in *History and Philosophy of Modern Mathematics*, W. Asprey and P. Kitcher eds., University of Minnesota Press, 1988, www.mcps.umn.edu/philosophy/11_15Aspray.pdf, accessed on October 2011.
- [A2005] William Aspray, On John von Neumann and his role in the development of the computer, *SIAM News* March 1, 2005, www.siam.org/news/news.php?id=37, accessed on October 2011.
- [ARD1986] Donald J. Albers, Constance Reid, and George B. Dantzig, An interview with George B. Dantzig: the father of linear programming, *The College Mathematics Journal* 17 (1986), 292–314.
- [B1951] G. W. Brown, Iterative solutions of games by fictitious play, in *Activity Analysis of Production and Allocation*, Cowles Commission for Research in Economics, Monograph No. 13, John Wiley & Sons, 1951, 374–376.
- [B2004] J. Bewersdorff, *Luck, Logic, and White Lies: The Mathematics of Games*, AK Peters, 2005.
- [BCMM1957] Roger R. Baldwin, Wilbert Cantey, Herbert Maisel, and James McDermott, *Playing Blackjack to Win: A New Strategy for the Game of 21*, M. Barrows and Company, 1957.
- [BG2002] Arthur T. Benjamin and Alan J. Goldman, Analysis of the N -card version of the game Le Her, *Journal of Optimization Theory and Applications*, 114 (2002) 695–704.
- [BS1979] Steven J. Brams and Philip D. Straffin Jr., Prisoner's Dilemma and professional sports drafts, *American Mathematical Monthly* 86 (1979) 80–88.
- [BT1999] Steven J. Brams and Alan D. Taylor, *The Win-Win Solution: Guaranteeing Fair Shares to Everybody*, W.W. Norton, 1999.
- [C1983] Vasek Chvatal, *Linear Programming*, W. H. Freeman & Co, 1983.
- [D1997] Morton D. Davis, *Game Theory: A Nontechnical Introduction*, Dover Publications, 1997.
- [D1961] Melvin Dresher, *Games of Strategy, Theory and Applications*, RAND corporation, 1961.
- [D2003] Ruchira S. Datta, Using computer algebra to find Nash equilibria, in *Proceedings of the 2003 international symposium on symbolic and algebraic computation*, ACM, 2003, 74–79.

- [D2013] Freeman Dyson, A walk through Johnny von Neumann's garden, *Notices of the AMS* 60 (2013).
- [Fn.d.] FairVote, Who picks the president? A report by FairVote, archive.fairvote.org/media/research/who_picks_president.pdf, accessed in February 2012.
- [GW1950] O. Gross and R. Wagner, A continuous Colonel Blotto game, RAND Research Memorandum 408, 1950.
- [GN1994] Ralph Gasser and Jürg Nievergelt, Es ist entschieden: Das Mühlespiel ist unentschieden, *Informatik Spektrum* 17(5) (1994) 314–317.
- [G1996] R. Gasser, Solving nine men's morris, *Computational Intelligence* 12 (1996) 24–41.
- [HL2010] Frederick S. Hillier and Gerald J. Lieberman, *Introduction to Operations Research*, McGraw-Hill Higher Education, 9th ed. 2010.
- [J1869] C. Jordan, Sur les assemblage de lignes, *Journal Reine Angew. Math.* 70 (1869) 185–190.
- [K1950] H. W. Kuhn, A simplified two-person poker, *Contributions to the Theory of Games* 1 (1950) 97–103.
- [K1953] H. W. Kuhn, Extensive games and the problem of information, *Contributions to the Theory of Games* 2 (1953) 193–216.
- [K1956] J. L. Kelly Jr., A new interpretation of information rate, *Bell System Technical Journal* 35 (1956) 917–926.
- [K1959] S. Karlin, *Mathematical Methods and Theory in Games, Programming and Economics*, Vol. I, Addison-Wesley, 1959.
- [K1961] Ephraim Kishon, Jewish Poker, 1961, www.ephraimkishon.de/my_favorite_stories.htm, last accessed September 2010.
- [K2003] Tinne Hoff Kjeldsen, New mathematical disciplines and research in the wake of world war II, in *Mathematics and War*, Bernhelm Booss-Bavnbek and Jens Høyrup, eds., Birkhäuser, 2003, 126–152.
- [K2004] H. W. Kuhn, Introduction, in the commemorative edition of John von Neumann & Oskar Morgenstern, *Theory of Games and Economic Behavior*, 2004, Princeton University Press, also at press.princeton.edu/chapters/i7802.pdf, accessed on October 2011.
- [K1995] V. Knoblauch, Geometric versions of finite games: prisoner's dilemma, entry deterrence and a cyclic majority paradox, *International Journal of Game Theory* 24 (1995) 165–177.
- [K2007] Marek Kaminski, Lecture 6, Sequential Games, in *Introduction to Game Theory*, webfiles.uci.edu/mkaminsk/www/gtp/06SequentialGames.pdf, accessed in November 2011.
- [Kn.d.] Levent Kockesen, Game Theory Lecture Notes, Lecture 5, undated web pages, home.ku.edu.tr/~lkockesen/teaching/uggame/lectnotes/uglect5.pdf, last accessed August 2011.
- [KC1971] David A. Kohler and R. Chandrasekaran, A class of sequential games, *Operations Research* 19 (1971) 270–277.
- [L2006] Steven M. LaValle, *Planning Algorithms*, Cambridge University Press, 2006.

- [LH1964] Carlton E. Lemke and J. T. Howson, Jr., Equilibrium points of bimatrix games, *SIAM Journal of Applied Mathematics* 12 (1964) 413–423.
- [MP1992] R. McKelvey and T. Palfrey, An experimental study of the centipede game, *Econometrica* 60 (1992) 803–836.
- [M1998] John Milnor, John Nash and *A Beautiful Mind*, *Notices of the AMS*, 45 (1998) 1329–1332.
- [N1998] Sylvia Nasar, *A Beautiful Mind*, Simon & Schuster, 1998.
- [N2001] Nobel Prize Committee, Interview with George A. Akerlof, A. Michael Spence, and Joseph E. Stiglitz, www.nobelprize.org/mediaplayer/index.php?id=502.
- [N2008] Nobel Prize Committee, Interview with Robert Aumann, www.nobelprize.org/nobel_prizes/economics/laureates/2005/aumann-interview.html.
- [P1993] William Poundstone, *Prisoner's Dilemma*, Random House, 1993.
- [P2005] William Poundstone, *Fortune's Formula*, Hill and Wang, 2005.
- [P2011] E. Prisner, Best response digraphs for two location games on graphs, in *Contributions to Game Theory and Management IV*, Leon A. Petrosjan and Nikolay A. Zenkevich, eds., St. Petersburg, (2011) 378–388.
- [PF1654] English translation (by Vera Sanford and Maxine Merrington) of the letters between Fermat and Pascal, www.york.ac.uk/depts/maths/histstat/pascal.pdf, accessed in December 2011.
- [PP2005] Knowing when to fold 'em, Q & A with Palisades' poker champ Chris Ferguson, *Palisadian Post*, 2005, www.palisadespost.com/content/index.cfm?Story_ID=1369.
- [R1951] J. Robinson, An iterative method for solving a game, *Annals of Mathematics* 54 (1951) 296–301.
- [R1981] R. Rosenthal, Games of perfect information, predatory pricing, and the chain store, *Journal of Economic Theory* 25 (1981) 92–100.
- [RC2004] Allan Rossman and Beth Chance, Conditional probability activity: AIDS testing, cs.zju.edu.cn/people/yedeshi/probability08/aidstesting.doc, accessed in September 2010.
- [S1971] Martin Shubik, The dollar auction game: A paradox in noncooperative behavior and escalation, *The Journal of Conflict Resolution* 15 (1971) 109–111.
- [S1960] Thomas C. Schelling, *The Strategy of Conflict*, Harvard University Press, 1960.
- [SBBKMLLS2007] Jonathan Schaeffer, Neil Burch, Yngvi Björnsson, Akihiro Kishimoto, Martin Müller, Robert Lake, Paul Lu, and Steve Sutphen, Checkers is solved, *Science* 317 (2007) 1518–1522.
- [S2009] Branislav L. Slantchev, Game theory: dominance, Nash equilibrium, symmetry, 2009, slantchev.ucsd.edu/courses/gt/04-strategic-form.pdf, accessed January 2012.
- [SW2001] U. Schwalbe and P. Walker, Zermelo and the early history of game theory, *Games and Economic Behavior* 34 (2001) 123–137.
- [T1962] Edward O. Thorp, *Beat the Dealer*, Blaisdell Pub. Co., 1962.

- [VNM1944] John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, 1953.
- [VS2007] Bernhard von Stengel, Lecture Notes on Game Theory, www.maths.lse.ac.uk/Courses/MA301/lectnotes.pdf 2007
- [Wn.d.] Dollar Auctions, a Wikipedia article at en.wikipedia.org/wiki/Dollar_auction.
- [Z1968] B. Zelinka, Medians and peripherians of trees, *Archivum Mathematicum* 4 (1968) 87–95.
- [Z1913] E. Zermelo, Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels, *Proceedings of the Fifth International Congress of Mathematicians* 2 (1913) 501–504.
- [ZP1993] Uri Zwick and Michael S. Paterson, The memory game, *Theoretical Computer Science* 110 (1993) 169–196.

Index

- 1-CANDIDATE QUIZ SHOW, 123
- 2 FINGERS, 56
- 2-AUCTION, 69
- 3-AUCTION, 69
- 5 CARDINALS, 45
- 5 ENVELOPES, 124
- 5 KNIGHTS, 12
- 5 PIRATES, 68
- 5 PIRATES with random selection, 126
- 8 ROUNDS BLACK OR WHITE, 147

- A Beautiful Mind, 212
- ABABA, 112
- ABABAB, 70
- ABBAAB, 70
- ABBABA, 70, 112
- action, 1
- acyclic digraph, 57, 160
- adjacent, 25, 34
- Akerlof, George A., 278
- alternatives, 160
- altruism, 60
- arc, 13
- arcs, 54
- Aumann, Robert, 157, 278

- backward induction
 - for perfect information and no randomness, 59
 - for perfect information and randomness, 122
- Banzhaf-Penrose Index, 52
- BARGAINING, 65
- BATTLE OF THE SEXES, 10
- behavioral strategy, 231
- best response, 8
- best response digraph, 13
- Borel, Emile, 63
- bottom-up analysis, 75
- branch subgame, 164
- Brown's fictitious play, 200
- Brown, George W., 156

- C&C, 113
- Card counting, 135
- cardinal scale, 90
- CENTIPEDE GAME, 60
- CHICKEN, 81
- CHOCOLATES IN THE BOX, 68
- COLONEL BLOTTO, 157
- complete information, 2
- conditional probability, 91
- cooperative game, 2
- core, 256
- CRASH ROCK-SCISSORS-PAPER, 201
- CREATING HOLES, 66
- cycle, 33
- cyclic digraph, 57

- Dantzig, George, 156, 204, 213
- digraph
 - acyclic, 160
 - probability, 93
- DMA6* SOCCER, 109
- DMA7* SOCCER, 109
- Dresher, Melvin, 158
- DUEL, 149, 152
 - two-bullet, 149, 152

- edge, 34
- ELECTION, 50, 157
- empirical probability, 87
- end positions, 53
- end vertex, 33, 160
- envy-free, 111
- equilibrium
 - Nash, 10
- event, 86
- Excel, 42
 - cells, 42
 - input cells, 42
 - output cells, 42
- experiments, 86

- extensive form, 54, 159
- fair, 80
- Ferguson, Chris, 138
- fictitious play, 200
- FINAL EXAM, 98
- final vertex, 160
- first mover advantage, 80
- Flood, Merrill, 158
- gain digraph, 72
- gain value, 72
- Gale, David, 279
- game, 1
 - finite, 56
 - symmetric, 6
 - zero-sum, 5
- game digraph, 54, 159
- game tree, 55, 160
- games
 - 1-CANDIDATE QUIZ SHOW, 123
 - 2 FINGERS, 56
 - 2-AUCTION, 69
 - 3-AUCTION, 69
 - 5 Cardinals, 45
 - 5 ENVELOPES, 124
 - 5 Knights, 12
 - 5 PIRATES, 68
 - 5 PIRATES with random selection, 126
 - 8 ROUNDS BLACK OR WHITE, 147
 - ABABA, 112
 - ABABAB, 112
 - ABBABA, 112
 - BARGAINING, 65
 - C&C, 113
 - CENTIPEDE GAME, 60
 - CHICKEN, 81
 - CHOCOLATES IN THE BOX, 68
 - COLONEL BLOTTO, 157
 - CRASH ROCK-SCISSORS-PAPER, 201
 - CREATING HOLES, 66
 - DMA6* SOCCER, 109
 - DMA7* SOCCER, 109
 - DUEL, 149
 - ELECTION, 50, 157
 - FINAL EXAM, 98
 - iterated COLONEL BLOTTO, 211
 - JOB INTERVIEWS, 127
 - LISA'S GAME, 69
 - MATCHING CHAIRS, 65
 - MINI BLACKJACK, 142
 - MINI BLACKJACK(a, b), 139
 - NIM(6), 53
 - OH-NO, 124
 - OH-NO6, 128
 - Polyomino REC THE SQUARE with Randomness, 125
 - Polyomino REC THE SQUARE with randomness, 128
 - PRISONER'S DILEMMA, 158
 - R&E2, 113
 - R&E4, 113
 - RANDOM NIM(n, p), 121
 - REC THE SQUARE, 66
 - ROCK-SCISSORS-PAPER, 81
 - ROCK-SCISSORS-PAPER2, 202
 - SEANATE RACE II, 54
 - SENATE RACE, 96
 - SEQUENTIAL LEGISLATORS VOTE, 57
 - SHUBIK AUCTION, 77, 169
 - Simultaneous Quatro-Uno, 44
 - symmetric, 80
 - TAKE SOME, 66, 69
 - TENNIS:, 97
 - ULTIMATUM GAME, 60
 - WHO'S NEXT, 67
 - x ROUNDS RED, GREEN, OR BLUE, 148
- graph, 34
- greedy strategy, 70
- guaranteed payoffs, 62
- Harsanyi, John, 181, 277
- Hurwicz, Leonid, 278
- IESD, 7
- IEWD, 8
- inefficiency, 111
- inequity, 111
- information set, 159
- interval scale, 90
- isomorphic, 157
- iterated COLONEL BLOTTO, 211
- JOB INTERVIEWS, 127
- Julia Robinson's theorem, 203

- Kelly criterion, 136, 137
- Kelly Jr., J. L., 135
- Kimmel, Emmanuel, 135
- Kuhn Poker, 186
- Kuhn, H.W., 185
- Kuhn, Harold William, 279
- law of large numbers, 87
- Le Her
 - Le Her*(S,r), 178
- linear program, 203
- LISA'S GAME, 69
- Marcum, Jess, 135
- Maskin, Eric S., 278
- MATCHING CHAIRS, 65
- maximal outerplanar graph, 33
- maximin move, 6
- maximin strategy, 6
- maximin value, 6
- median of a graph, 30
- Min Sat, 112
- MINI BLACKJACK, 142
- MINI BLACKJACK(a, b), 139
- Mirrlees, James A., 278
- mixed maximin strategy, 202
- mixed Nash equilibrium, 203
- mixed security level, 202
- mixed strategy, 199
- MOP, 33
- Morgenstern, Oskar, 185
- move, 1
- moves, 160
- Myerson, Roger B., 278
- Nash equilibrium
 - mixed, 203
 - Pareto-optimal, 11
 - pure, 10
 - subgame perfect, 181
- Nash's Theorem, 204
- Nash, John, 13, 135, 156, 181, 212, 277
- neighbor, 34
- net odds, 136
- NIM(6), 53
- Nobel Memorial Prize in Economics, 13, 157
- Nobel, Alfred, 277
- nominal scale, 90
- normal form, 4, 5, 177
- OH-NO, 124
- OH-NO6, 128
- options, 160
- ordinal scale, 90
- outcome, 4
- outcomes, 1
- Pareto-dominated, 11, 111
- Pareto-optimal, 11
- payoff, 1
- perfect information, 2
- perfect recall, 233
- play, 1
- Polyomino REC THE SQUARE with Randomness, 125
- Polyomino REC THE SQUARE with randomness, 128
- position, 1, 53
- PRISONER'S DILEMMA, 10, 158
- probability, 87
 - empirical, 87
 - theoretical, 87
- probability digraph, 93
- probability tree, 91
- projects
 - 5 Cardinals, 45
 - 5 ENVELOPES, 127
 - 5 Knights, 12
 - JOB INTERVIEWS, 127
 - OH-NO, 128
 - Polyomino REC THE SQUARE with randomness, 128
 - Restaurant Location Games, 44
 - Simultaneous Quatro-Uno, 44
- proof, 15
- pure Nash equilibrium, 10
- pure strategy, 176
- R&E2, 113
- R&E4, 113
- random moves, 2
- RANDOM NIM(n, p), 121
- random variable, 89
- random vertex, 121
- ratio scale, 90
- REC THE SQUARE, 66

- reduced pure strategy, 177
- relative frequency, 87
- relative values, 111
- ROCK-SCISSORS-PAPER, 81
- ROCK-SCISSORS-PAPER2, 202
- Roth, Alvin E., 278
- satisfaction, 111
- scale
 - cardinal, 90
 - interval, 90
 - nominal, 90
 - ordinal, 90
 - ratio, 90
- Schelling, Thomas, 157, 278
- second mover advantage, 80
- security level, 6, 62
- security level analysis, 62
- security strategy, 62
- Selten, Reinhard, 181, 277
- SENATE RACE, 96
- SENATE RACE II, 54
- sequential game, 2, 53
- SEQUENTIAL LEGISLATORS VOTE, 57
- set position, 72
- Shapley value, 256
- Shapley, Lloyd, 278
- Shapley, Lloyd Stowell, 135, 156
- SHUBIK AUCTION, 77, 169
- Shubik, Martin, 135, 279
- simple event, 86
- simultaneous game, 2
- Simultaneous Quatro-Uno, 44
- Spence, A. Michael, 278
- spreadsheet program, 42
- stable coalition, 255
- standard deviation, 115
- start position, 53
- start vertex, 160
- Stiglitz, Joseph E., 278
- strategic form, 177
- strategy, 1
 - pure, 176
 - reduced, 177
- strict domination, 7
- subgame, 164
- subgame perfect, 181
- support, 200
- swing voter, 52
- symmetric games, 80
- TAKE SOME, 66, 69
- TENNIS, 97
- terminal vertex, 160
- theorem, 15
- theoretical probability, 87
- Thorp, Edward O., 135
- ULTIMATUM GAME, 60
- unique response, 84
- Uston, Ken, 138
- utilities, 60
- utility function, 136
- value of a two-player zero-sum game, 202
- variance, 116
- vertex, 13, 34
- vertices, 54
- Vickrey, William, 278
- VNM Poker, 185
- von Neumann's Theorem, 203
- von Neumann, John, 63, 156, 185, 212
- weak domination, 8
- WHO'S NEXT, 67
- Williams, John David, 157
- winning coalition, 52
- x ROUNDS RED, GREEN, OR BLUE, 148
- Zermelo's Theorem, 59
- Zermelo, Ernst, 59
- zero-sum, 5
- zero-sum game, 2